

Delayed Verification Mechanism for Dynamic Implementation*

Olga Gorelkina[†]

June 25, 2014

Abstract

This paper introduces a virtually efficient mechanism in a setting with sequentially arriving agents who hold informative signals about future types. To reveal the information the principal organises betting on future type reports. An agent's betting reward depends on how accurately the prior updated on his report predicts the type reports observed in the following period. The mechanism satisfies participation constraints and generates no deficit after any reported history.

JEL: D61, D71, D82, D84

Keywords: Dynamic mechanisms, scoring rule, Bayesian learning, conjugate priors.

1 Introduction

This study is motivated by a number of practical applications of dynamic mechanism design, where the principal lacks information about the future type distri-

*This study has been funded by the Max Planck Society and Toulouse School of Economics. Further acknowledgments to be added.

[†]Max Planck Institute, Bonn.

butions. In industrial regulation, pollution control or auctions such information is necessary to evaluate the intertemporal trade-offs, but often unavailable to the principal. Due to Mezzetti (2004), we know that full revelation of future-related information is achieved by delaying incentive payments until the resolution of all uncertainty. However in practice there may be obstacles to delaying payments in mechanisms that run for a long time. Gershkov and Moldovanu (2009) are the first to study the problem of principal’s learning under this timing constraint.¹ They look at a setting where the present agents’ payoff types are informative of the future, therefore the principal can make inferences from the current type reports. In this paper we take a different approach and suppose that the information about the future is independent of agents’ own preferences. We show by means of an example that revelation is non-trivial also in this case. By constructing a mechanism with verification we prove that dynamic efficiency can virtually be achieved, if transfers are delayed just until the next cohort’s arrival.

Formally we consider a discrete time, finite horizon setting and agents who arrive sequentially in cohorts and live for two periods.² A member of cohort t derives utility from allocation at t ; in period $t + 1$ he receives monetary transfers, and his total utility is the sum of both components. As standard in this literature, the agent has parametrizable private information composing his *type*. The type here has two independent components: (i) a one-dimensional preference parameter, referred to as the *payoff type*, and (ii) a multidimensional signal drawn from the future payoff type distribution, referred to as the *hyperbelief type*.³ The

¹Dynamic populations have also been extensively studied in the literature often referred to as online mechanism design: see the seminal works by Lavi and Nisan (2004) and Parkes and Singh (2003). However their approach is non-Bayesian, and therefore the problem of principal’s learning does not emerge.

²The population is thus dynamic with static private information, similarly to Lavi and Nisan (2004) and Parkes and Singh (2003). Another strand of literature studies persistent population with dynamic types: Athey and Segal (2013), Bergemann and Välimäki (2010). Cavallo et al. (2009) study the general setting with arbitrary dynamic population and dynamic types.

³The term *hyperbelief* comes from *hyperdistribution*, a term borrowed from the theory of conjugate priors used in the present analysis (see Raiffa and Schleifer (1961)). It should not be confused with higher-order belief as it describes the probabilistic view of the future environment, and not of other agents’ beliefs.

payoff type pins down the agents' own utility function, while the hyperbelief type reflects his information about the future.

The parametric class of payoff-type distributions is fixed, however the exact distribution is unknown. Let p_t be the probability distribution with an unknown parameter α_t and Φ_t be a class of (hyper-)distributions over the possible values of α_t . A hyperbelief type is then draw from a future p_t which leads to an updated $\phi_t \in \Phi_t$. We assume that the class Φ_t is *conjugate* to p_t , meaning that all updated hyperdistributions belong to the same class. Conjugate classes have been found for most standard distributions (Raiffa and Schleifer, 1961).

The delayed verification mechanism introduced here features a two-part transfer that work towards elicitation of both components of type.⁴ Firstly, to elicit the payoff type we introduce a Vickrey-Clarke-Groves transfer (1961; 1971; 1973). The VCG transfer equals the externality imposed by the agent's payoff-type report on the current and the future generations, where the future welfare is taken in expectation conditional on the entire history of reported hyperbeliefs. To this end, we follow the existing literature on dynamic implementation, such as Bergemann and Välimäki (2010); Pavan et al. (2010).⁵

The second part of transfer, novel to dynamic implementation, is the scoring reward used for the verification of signals. It is a function of agent's report of hyperbelief type (signal) and the next cohort's type reports. More precisely, the agent receives a (negative) transfer equal to the log-likelihood of the profile of type reports in the following period, where the likelihood is evaluated according to the distribution updated conditionally on *his* report. Reporting hyperbelief to the mechanism is thus similar to placing a bet on the next period type reports. Given that signals and payoff-types are drawn independently, matching the empirical distribution of payoff types is equivalent to matching the empirical distri-

⁴Parkes and Singh (2003) use the term *delayed mechanism* to describe a mechanism, where all the payments are made in the last period (see Friedman and Parkes, 2003). Note that this is not the case here: the transfers are made with only one period delay.

⁵In a setting with persistent population and dynamic types, Athey and Segal (2013) design a dynamic extension of the expected externality mechanism d'Aspremont et al. (2004) satisfying the exact budget balance, unlike the VCG.

bution of signals. The expected value of such scoring transfer is maximized when the hyperbelief types are reported truthfully. In static implementation mutual verification of reports has been used in the literature stemming from Crémer and McLean, 1985.⁶

The mechanism with delayed verification features linear combination of the VCG and scoring transfers. This yields *virtual* efficiency in the following sense. For any given value $\varepsilon > 0$, the principal can scale up the scoring transfer such that truthful revelation of payoff *and* hyperbelief types is an ε -equilibrium. In other words, no agent can gain at least ε in deviation from truth-telling. Moreover all small deviations are in the neighborhood of truth-telling, thus the present concept is stronger than ε -equilibrium. The impossibility of exact implementation in this setting is due to the fact that belief report marginally interferes with the externality payment.

Finally, we show that the budget of bets can be exactly balanced without change to incentives. Assume that each cohort includes at least two participants, and define an arbitrary derangement of the set of players in the cohort. The derangement defines cycles of payments between the agents in a cohort, such that the agents pay their betting rewards to each other, and the exact budget balance in bets is achieved in every period. The resulting transfer, the *balanced scoring* transfer, satisfies the individual rationality constraint. The intuition for this is that the ex ante distribution of signals is the same for all agents; therefore, the expected difference in betting rewards is zero. It follows that the entire *balanced* mechanism that comprises the VCG and the balanced scoring transfers requires no external funding at any period of time (generates no deficit after any history) and satisfies the participation constraint.

The rest of the paper is organized as follows. Section 2 considers the allocation of pollution permits in a simple setting with two periods. In this example the

⁶Crémer and McLean (1985, 1988) were the first to design a mechanism with mutual report verification in a static setting. Their mechanism punishes for reports that appear contradictory given the known correlation between types, and thus reveals the types perfectly. McAfee and Reny (1992) provide an extension of the Crémer-McLean mechanism to continuous type spaces.

comparison of reports within one period (as in Crémer and McLean, 1985) fails to produce truth-telling as a unique and undominated equilibrium; therefore delayed verification is introduced. Section 3 presents a more general setting with an arbitrary, although finite number of periods. Section 4 presents the delayed verification mechanism. Two subsections (4.1 and 4.2) study the two components of the delayed-verification transfer: the VCG and the scoring payment. In 4.1 we show that the VCG transfer induces the revelation of payoff-type, if the current and the future hyperbelief-types are truthfully reported (Lemma 1). In 4.2 we show that the scoring transfer alone ensures that hyperbelief types are truthfully reported, if the future payoff-type reports are truthful (Lemma 2). Finally in 4.3 we show that both transfers can be combined to achieve truthful revelation of the entire type in an ε -equilibrium (Proposition 1). A mechanism with exactly balanced scoring transfer can be constructed to satisfy the ex ante participation constraint and no deficit in transfers after any history of reports (Proposition 2). A summary of notation is given in the Appendix.

2 Illustration: Two-period pollution control.

Delayed verification in a nonrepeated setup is illustrated in a simple model of pollution control. Suppose an area accommodates n firms that use a hazardous input Y for their production in period 1. A firm's cost of production decreases in its usage of Y :

$$C(y_i; \theta_i) = \frac{1}{\theta_i y_i},$$

where $y_i > 0$ is the amount of Y used by firm $i = 1, 2, \dots, n$ and θ_i is a technology parameter privately known to the firm. The regulator's task is to allocate usage permits to firms (y_1, y_2, \dots, y_n) such that the total cost of production and environmental damage are minimized.

The environmental damage in period 2 is linear in the total use of Y and equals:

$$\theta_d \sum_i y_i,$$

where θ_d is the damage per unit of Y used (θ_d is the *type* of environment at $t = 2$). We assume that θ_d is a Bernoulli random variable: it takes value $D > 0$ with probability α or value 0 with probability $1 - \alpha$. Denote $p(\theta_d; \alpha)$ the respective probability distribution function:

$$p(\theta_d; \alpha) = \begin{cases} \alpha, & \theta_d = D \\ 1 - \alpha, & \theta_d = 0 \end{cases}.$$

Contrary to the common assumption in mechanism design, we consider the high damage probability α unknown. Each firm i observes an informative signal,

$$x_i \in \{0, D\},$$

drawn independently from distribution $p(\theta_d; \alpha)$. Note that any such draw is informative of α . Let $X = (x_1, x_2, \dots, x_n)$ denote the vector of observed signals. X represents all the information available at $t = 1$ about the damage at $t = 2$.

The efficient allocation of permits in period 1 minimizes the expected total cost to all agents (we assume no discounting between periods), provided the information X :

$$\min_{(y_1, y_2, \dots, y_n)} \int_{[0,1]} \left(\sum_i \frac{1}{\theta_i y_i} + \alpha D \sum_i y_i \right) \phi(\alpha | X) d\alpha, \quad (1)$$

where the probability measure $\phi(\alpha | X)$ is the Bayes-updated hyperprior $\phi^0(\alpha)$ given X :

$$\phi(\alpha | X) = \frac{\Pr[X|\alpha] \phi^0(\alpha)}{\int_{[0,1]} \Pr[X|\tilde{\alpha}] \phi^0(\tilde{\alpha}) d\tilde{\alpha}}, \quad (2)$$




Figure 1: *The principal's initial hyperbelief* $\phi^0(\alpha) = \beta(\alpha; 1, 1)$, *the data* $X = (\theta_d^L, \theta_d^L, \theta_d^H)$, *and the updated hyperbelief* $\phi(\alpha|X) = \beta(\alpha; h + 1, l + 1) = \frac{\alpha^h(1-\alpha)^l}{B(2,3)} = \frac{\alpha(1-\alpha)^2}{B(2,3)}$.

where $\Pr[X|\alpha] = \prod_i p(x_i|\alpha)$.

Suppose that the prior $\phi^0(\alpha)$ is uniform over the interval $[0, 1]$ of the possible values of α . Since the uniform distribution belongs to the class of Beta distributions, conjugate to the Bernoulli class, the update $\phi(\alpha|X)$ is also a Beta distribution.⁷ If $h = \#\{x_i = D\}$ denotes the number of firms with a high signal, and $l = \#\{x_i = 0\}$ the number of firms with a low signal, then the updated probability (2) becomes:

$$\phi(\alpha|X) = \frac{\alpha^h (1 - \alpha)^l \beta(\alpha; 1, 1)}{\int_{[0,1]} \tilde{\alpha}^h (1 - \tilde{\alpha})^l \beta(\tilde{\alpha}; 1, 1) d\tilde{\alpha}} = \beta(\alpha; 1 + h, 1 + l) \quad (3)$$

Note that h and l are sufficient statistics for data X . Given the updated prior, the efficient allocation of permits that solves the minimization problem (1) is the following:

$$f_i^*(\theta, X) = \sqrt{\frac{2 + h + l}{\theta_i D (1 + h)}}, \quad (4)$$

for all $i = 1, 2, \dots, n$.

The difficulty in implementing the efficient allocation f^* is that there is no guar-

⁷The uniform distribution, or $\beta(\alpha; 1, 1)$, has the maximal entropy within the class of Beta distributions, and thus is the least informed prior within that class.

antee that firms will reveal the signals about the future truthfully (i.e. the numbers h and l may not represent the actual number of high and low signals received). This implies that the mechanism has to be incentive-compatible not only with regard to the elicitation of cost parameters, but also of the signals. As a possible solution to this problem we first discuss an application of the idea of Crémer and McLean (1985) as the mutual verification of firms' reports in the first period (immediately). In the second step we show that welfare is improved by delaying verification to the second period when the damage realizes. In Sections 3 and 4 we extend the result to a repeated setting.

Immediate Verification Crémer and McLean (1985) induce truthfulness by rewarding similarity and punishing divergence in the agents' reports. Consider a direct mechanism, where each firm reports its technology parameter θ_i and the signal about future damage x_i . The principal updates her prior ϕ^0 on the firms' information $X = (x_i)_{i \in N}$ and assigns the efficient allocation f^* . With logarithmic scoring, we have the following two-part transfer to firm i :

$$\tau_i(\hat{\theta}_i, \hat{X}) = \underbrace{-\mathbb{E}[\theta_d | \hat{X}] \times f_i^*(\hat{\theta}_i, \hat{X})}_{\text{externality}} + \lambda \times \underbrace{\ln \prod_{j \neq i} \Pr[\hat{x}_j | \hat{x}_i]}_{\text{scoring}}, \quad (5)$$

where $\hat{X} = (\hat{x}_i)_{i \in N}$ denotes the reported values, and $\Pr[\hat{x}_j | \hat{x}_i]$ is calculated as follows. If $\hat{x}_j = D$, $\Pr[\hat{x}_j | \hat{x}_i] = \int_{[0,1]} \alpha \phi(\alpha | \hat{x}_i) d\alpha$ and if $\hat{x}_j = 0$, $\Pr[\hat{x}_j | \hat{x}_i] = \int_{[0,1]} (1 - \alpha) \phi(\alpha | \hat{x}_i) d\alpha$.

The transfer consists of two parts that provide distinct incentives. The *externality* transfer induces the truthful revelation of the technology parameter, whereas the *scoring* transfer rewards similarity in reports.

Claim 1. For any $\lambda > 0$ reporting no damage ($D = 0$) irrespectively of the true signal is an equilibrium in the immediate verification mechanism. This equilibrium Pareto dominates the truthful revelation of signals.

(See proof in the Appendix) The claim says that the firms will be prone to report low damages $\hat{x}_i = 0$ even if their true signal is high $\hat{x}_i = D$. As a result of this manipulation, the regulator would underestimate the probability of a high damage and assign larger permits y_i in the first period, such that the firms will save production costs creating higher-than-efficient pollution.

Delayed Verification The mechanism with delayed verification provides incentives for truthful revelation of signals. The change is made to the scoring transfer. Now, instead of rewarding similarity in firms' reports, the mechanism rewards similarity between reports and the realized damage. The mechanism assigns the following transfers after the environmental uncertainty resolves. If the realized damage is high, $\theta_d = D$, then firm i will receive the following two-part transfer:

$$\tau_i(\hat{\theta}_i, \hat{X}; D) = \underbrace{-D \times f_i^*(\hat{\theta}_i, \hat{X})}_{\text{externality}} + \lambda \times \underbrace{\ln \int_{[0,1]} \alpha \phi(\alpha|\hat{x}_i) d\alpha}_{\text{scoring}}, \quad (6)$$

If the realized damage is low, $\theta_d = 0$, the transfer amounts to:

$$\tau_i(\hat{\theta}_i, \hat{X}; 0) = \underbrace{0 \times f_i^*(\hat{\theta}_i, \hat{X})}_{\text{externality}} + \lambda \times \underbrace{\ln \int_{[0,1]} (1 - \alpha) \phi(\alpha|\hat{x}_i) d\alpha}_{\text{scoring}}. \quad (7)$$

The scoring transfer is the “betting reward” that the firm gets for correctly predicting the realized damage. Observe that $\int_{[0,1]} \alpha \phi(\alpha|\hat{x}_i) d\alpha$ and $\int_{[0,1]} (1 - \alpha) \phi(\alpha|\hat{x}_i) d\alpha$ are the conditional likelihoods of, respectively, D and 0 given i 's reported signal x_i . Lemma 2 in the main text states, for a more general case, that the maximization of the scoring transfer induces truthful report of signals.

Claim 2. Maximization of the delayed verification scoring transfer yields truthful revelation of signal as a strict optimum.

When signals are reported truthfully, the externality (VCG) transfer represents the actual change in the social welfare due to the firm’s use of hazardous input Y . Since the firm pays the amount of the environmental damage its production has caused, its objective is aligned with the principal’s program of total welfare maximization, and therefore it is optimal to report the true cost parameter θ_i . Observe that hyperbelief report \hat{x}_i also enters the VCG transfer through the update of principal’s prior $\phi(\alpha|X)$ used to estimate the future externality. In settings with continuous types, this leads to small deviations from truth-telling. In an attempt to marginally change the prior the agent slightly misreports. Significant deviations can be precluded by scaling up the scoring transfer. The proof of Proposition 1 demonstrates that for any given level of precision one can find $\lambda \in \mathbb{R}$ to scale up the scoring transfer and preclude deviations from truth-telling. This result implies the virtual implementation of the efficient rule f^* .

Next we extend the idea of delayed verification to a finitely repeated setting. In that framework, the example of pollution control corresponds to the last two stages of the mechanism.

3 The General Model

Consider a dynamic system with finite sequence of opening periods $\{1, 2, \dots, T\} \equiv \mathbb{T}$, where $T \geq 2$. The set of participants entering at period t is called *cohort t* and denoted by N_t , each cohort includes at least two members. Let $N^t = \bigcup_{s \leq t} N_s$ denote the set of participants arriving at time t or earlier.⁸

The principal chooses an allocation from a time-invariant space \mathbb{Y} of alternatives; \mathbb{Y} is a compact subset of a metric space. $y_t \in \mathbb{Y}$ is the decision taken by the principal at $t \in \mathbb{T}$. The history of allocations up to t is denoted y^t .

⁸The generalization to setting where the agents live for more than two periods is straightforward, as long as the payoff type is persistent.

Utilities Agent $i \in N_t$ derives utility u_i from the allocation history y^t at t and receives monetary transfer at $t + 1$. The agent's utility function is known up to parameter $\theta_i \in \Theta_t$, where θ_i (also denoted x_{i0}) is i 's private information referred to as i 's *payoff type*.

$$u_i : \mathbb{Y}^t \times \Theta_t \rightarrow \mathbb{R}_+. \quad (8)$$

The utility functions u_i are strictly concave and Lipschitz-continuous in y_t for all $t, i \in N_t$. Additionally, I assume that non-participation yields zero utility to any agent $i \in N$.

Note that the utility functions are "observationally measurable"⁹ in the sense that the utility does not depend on any variables unobserved by the agent. In particular, the utility is invariant in the future allocations and other agents' types.

Types The payoff-type $\theta_i \equiv x_{i0}$ of agent $i \in N_t$ is a random draw from a compact set $\Theta_t \subseteq \mathbb{R}$, $t \in \mathbb{T}$. The corresponding probability distribution function $p_t(\cdot; \alpha_t)$ over Θ_t belongs to a given parametric class \mathcal{P}_t ; parameters α_t are serially uncorrelated. The class of priors conjugate to \mathcal{P}_t is denoted $\Phi_t \subset \Delta(\mathcal{A}_t)$, where \mathcal{A}_t is the set of possible values of parameter α_t . Let $\phi_t^0 \in \Phi_t$ denote the initial prior, common to agents N_1 and the principal. $\phi^0 = \prod_{t \in \mathbb{T}} \phi_t^0$ is the joint prior over the space of possible values $(\alpha_1, \alpha_2, \dots, \alpha_T)$ before the start of the game:

$$\phi^0 : \prod_{t \in \mathbb{T}} \mathcal{A}_t \rightarrow \mathbb{R}_+. \quad (9)$$

To distinguish the elements of Φ_t from the elements of \mathcal{P}_t we refer to $\phi_t \in \Phi_t$ as *hyperbeliefs* and to its parameters as *hyperparameters*. In the example of Section 2, \mathcal{P}_t is the Bernoulli class, Φ_t is the conjugate Beta class and $(1, 1)$ are the initial hyperparameters (*Beta*(1, 1) is the uniform distribution over the interval $[0, 1]$).

⁹As in Athey and Segal (2013)

Information Next to his own payoff type θ_i the agent holds private information about the future payoff-type distributions ϕ_t . For every t and $i \in N_t$ we define x_{ik} as i 's signal about cohort N_{t+k} . x_{ik} is a random draw from distribution $p_{t+k}(\cdot; \alpha_{t+k})$. Agent i 's entire private information, or *type*, is summarized in the vector x_i :

$$x_i = \left(\begin{array}{c} x_{i0} \equiv \theta_i \\ x_{i1} \\ \vdots \\ x_{iT-t} \end{array} \right) \left. \vphantom{\begin{array}{c} x_{i0} \equiv \theta_i \\ x_{i1} \\ \vdots \\ x_{iT-t} \end{array}} \right\} \begin{array}{l} \text{payoff-type} \\ \text{hyperbelief-type} \end{array} \quad (10)$$

where for all $i \in N_t$:

$$x_i \in \prod_{k=0}^{T-t} \Theta_{t+k}. \quad (11)$$

Let $X_t = (x_i)_{i \in N_t}$ denote the (matrix of) signals of cohort N_t and $X^t = (X_1, X_2, \dots, X_t)$ the history of signals up to period t . In the following, we refer to X^t as the *information available at t* , even though X^t cannot be observed entirely by any single agent at t . X^t comprises X^{t-1} , the information that is public at t ,¹⁰ as well as $|N_t|$ pieces of private information.

Sequential Updating The hyperbeliefs are updated upon the arrival of new information. In a given period $t \in \mathbb{T}$, types x_{i0} of the members of N_t and their signals x_{ik} about the future types are drawn. Given the hyperprior $\phi_s(\alpha_s | X^{t-1})$, $s > t$, from the previous period¹¹ and the new data X_t , the updated hyperprior is derived by the Bayes rule as follows:¹²

¹⁰In Lemma 3 we argue that the history of reports should be made public.

¹¹If $t = 0$ we have $\phi_s(\alpha_s | \emptyset) = \phi^0(\alpha_s)$ for all $s > 0$.

¹²The first equality in (12) holds due to the mutual independence of parameters.

$$\phi_s(\alpha_s|X^t) = \phi_s(\alpha_s|X_{t,s-t}, X^{t-1}) = \frac{\Pr[X_{t,s-t}|\alpha_s, X^{t-1}] \phi_s(\alpha_s|X^{t-1})}{\int_{\mathcal{A}_s} \Pr[X_{t,s-t}|\tilde{\alpha}_s, X^{t-1}] \phi_s(\tilde{\alpha}_s|X^{t-1}) d\tilde{\alpha}_s}, \quad (12)$$

where

$$X_{t,s-t} = (x_{1,s-t}, x_{2,s-t}, \dots, x_{|N_t|,s-t}) \quad (13)$$

is the profile of the signals of received by cohort t about the payoff-types at s , and

$$\Pr[X_{t,s-t}|\alpha_s, X^{t-1}] = \Pr[X_{t,s-t}|\alpha_s] = \prod_{i \in N_t} p_s(x_{i,s-t}; \alpha_s) \quad (14)$$

is the probability of such signal profile conditional on parameter value α_s .

It is easy to check that the update defined by Equation (12) is a probability measure. Moreover, by the conjugate prior property the updated hyperdistribution $\phi_s(\alpha_s|X^t)$ given by Equation (12) belongs to the same class as $\phi_s(\alpha_s|X^{t-1})$, for any $t \in \mathbb{T}$.

For $s \leq t$, the hyperprior is transferred from the previous period: $\phi_s(\alpha_s|X^t) = \phi_s(\alpha_s|X^s)$.¹³ Finally, the joint hyperprior writes:

$$\phi(\alpha|X^t) = \prod_{s \in \mathbb{T}} \phi_s(\alpha_s|X^t) \quad (15)$$

Efficiency The social welfare is defined as the sum¹⁴ of all the agents' utilities:

$$W(X^T, y^T) = \sum_{t \in \mathbb{T}} \sum_{i \in N_t} u_i(y^t, x_{i0}) \quad (16)$$

¹³The implementation problem at t does not require the hyperbelief over α_t to be updated.

¹⁴The discount factors δ^t can be subsumed in the utility functions $u_i, i \in N_t$.

A dynamic choice rule $f = (f_1, f_2, \dots, f_T)$ is a finite sequence of functions f_t mapping the information available up to t into the set of allocations \mathbb{Y} . A dynamic choice rule f is *dynamically efficient* if for all $t \in \mathbb{T}$ and X^t :

$$f_t(X^t) \in \arg \max_{y_t \in \mathbb{Y}} \{ \mathbb{E} [W(X^T, y^T) | X^t] \} \quad (17)$$

Observe that efficiency requires that expectation be conditional on the entire information available at t , including the information which is private.

We can reformulate the problem to derive a notion of efficiency that is more operational for our purpose of designing transfers. The choice rule f is dynamically efficient if it solves the *stochastic optimal control* problem, where the allocation y_t is the control variable and X^t, y^{t-1} is the state with a stochastic component. The law of motion of the stochastic component x_{ik} of state is given by:

$$x_{ik} \sim p_{t+k}(\cdot; \alpha_{t+k}) \quad (LM)$$

Provided the optimal control formalism, we can use the standard techniques to solve the dynamic problem (Bellman, 1966). Write the Bellman function as follows:

$$J_t(X^t, y^{t-1}) = \max_{y_t \in \mathbb{Y}} \left\{ \sum_{i \in N_t} u_i(y^t, x_{i0}) + \mathbb{E} [J_{t+1}(X^{t+1}, y^t) | X^t] \right\}, \quad (18)$$

subject to (LM) and the terminal condition $J_{T+1}(X^{T+1}, y^T) = 0$, further (TC). The Bellman function is interpreted as the maximal future value of the mechanism from time t on, given the past decisions and information. The value at t includes the known utilities of cohort N_t , and the future cohorts' utilities in expectation over their types.

By the Bellman principle we maximize (18) with respect to the control variable y_t , conditional on the information available at t , X^t . The only relevant uncer-

tainty at this stage is the uncertainty about X_{t+1} ; all the subsequent payoff-type uncertainty is contained in the Bellman value at $t + 1$.

The expected value of the Bellman function:¹⁵

$$\mathbb{E} [J_{t+1} (X^{t+1}, y^t) | X^t] = \int_{\mathcal{A}} \left(\int_{\mathbb{X}_{t+1}} J_{t+1} (X^{t+1}, y^t) \Pr [X_{t+1} | \tilde{\alpha}, X^t] d\mathbf{X}_{t+1} \right) \phi (\tilde{\alpha} | X^t) d\tilde{\alpha} \quad (19)$$

The following definition of efficiency is then equivalent to the one introduced previously.

Definition The choice rule f is dynamically efficient, if for all $t \in \mathbb{T}$ it solves the maximization problem in Equation (18).

Mechanisms A direct mechanism (f, τ) is an observationally-measurable mapping from types into allocation and transfers. We focus on mechanisms where transfers can be delayed by one period. That is, cohort N_t receives transfers at $t + 1$ (Figure 2). A mechanism *virtually implements* social choice function f if for any $\varepsilon > 0$ truth-telling is an ε -equilibrium, and any profitable deviation that yields payoff less than ε is in the neighborhood of truth-telling.¹⁶

Mechanism (f, τ) satisfies the *participation constraint* if for all $t \in \mathbb{T}$, $i \in N_t$ and X^{t-1} the following holds:

$$\mathbb{E} [u_i (f^t (X^t), x_{j0}) + \tau_i (X^{t+1}) | X^{t-1}] \geq 0 \quad (20)$$

The condition postulates that under “the veil of ignorance”, that is, before the agent observes his private information x_i , but after the observation of history

¹⁵Where $\Pr [X_{t+1} | \tilde{\alpha}, X^t] = \prod_{s=t+1, \dots, T} \Pr [X_{t+1, s-(t+1)} | \tilde{\alpha}_s] = \prod_{s=t+1, \dots, T} \prod_{i \in N_{t+1}} p_s (x_{i, s-(t+1)}; \alpha_s)$ is defined in Equation (14).

¹⁶A strategy profile is an ε -equilibrium, if there is no player and deviation that increases the player’s payoff by at least ε .

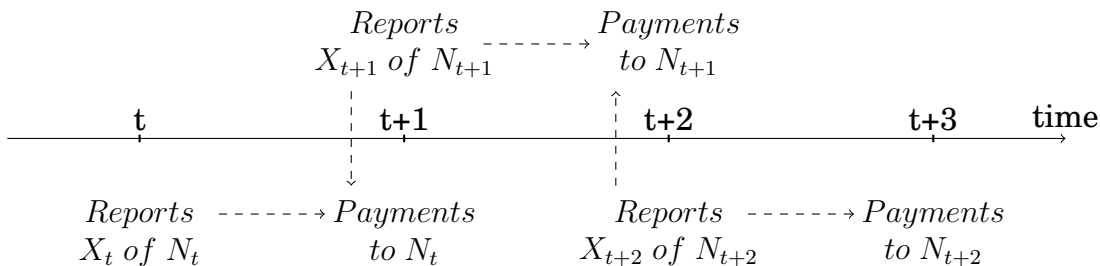


Figure 2: The timing of the mechanism.

X^{t-1} – the participation in the mechanism yield a higher payoff in expectation than abstention.¹⁷

Mechanism (f, τ) satisfies *no deficit*, if for all $t \in \mathbb{T}$, and all report histories X^{t+1} : $\sum_{i \in N_t} \tau_i(X^{t+1}) \leq 0$.

Note that this condition is stronger than the requirement of no deficit after the final round. It postulates that the payments made at every given point in time generate no deficit in the principal's budget.

4 The Delayed Verification Mechanism

The timing is described by the following iterations ($t = 1, 2, \dots, T - 1$).

- In period t agent $i \in N_t$ reports her information $x_i = (x_{i0}, x_{i1}, x_{i2}, \dots, x_{iT-t})$.
- The principal merges new data $X_t = (x_i)_{i \in N_t}$ with history X^{t-1} to obtain X^t , and updates his belief and implements allocation $f(X^t)$ at t .
- In period $t + 1$ agent $i \in N_t$ receives the transfer and leaves the mechanism.
Etc.

¹⁷If the participation condition 20 holds for agent $i \in N_t$, one can achieve that his payoff at $t + 1$ is positive as follows. At arrival in t , the agent buys a 0-interest bond from the principal and gets repayment at $t + 1$.

The *delayed verification mechanism* (f, τ) is a pair of efficient dynamic choice rule f , defined by Equation (18), and transfer system τ , defined as follows:

$$\tau_i(X^{t+1}) = \tau_i^{VCG}(X^t) + \lambda \tau_i^{Sc.}(X^{t+1}), \quad (21)$$

for all $t \in \mathbb{T}$ and $i \in N_t$, where the components are given by Equations (23) and (26), respectively.

The two-part transfer τ_i serves to induce truthful revelation of the entire vector of private information x_i and implement the efficient allocation. Essentially it formalizes the idea voting on preferences, but betting on beliefs (Hanson, 2013).

Next we discuss the construction of both parts of the transfer. Lemmas 1 and 2 state that the respective transfers yield truth-telling when applied separately to reveal types and hyperbeliefs; Lemma 3 is an auxiliary result on information disclosure by the principal. Proposition 1 states the virtual implementation result. Finally, we describe a way to balance the scoring budget and satisfy the ex-ante participation constraints, summarizing the result in Proposition 2. As before, all proofs are given in the Appendix.

4.1 VCG Transfer τ_i^{VCG}

The Vickrey-Clarke-Groves payment equals the externality, expressed in money, that the agent's report imposes on the present and future agents. This transfer aligns the incentives of every agent with the principal's objective, since the maximization of his own utility in sum with the VCG transfer is equivalent to total welfare maximization. Therefore truthful report of payoff type $\theta_i = x_{i0}$ is optimal.

To construct the VCG transfer in our environment, we first introduce a family of choice rules f_{-i} , $i \in N^T$, that are efficient with respect to a restricted player set N^T/i . That is, for a given i , f_{-i} maximizes the total welfare net of the utility of agent i :

$$W_{-i}(X^T, y^T) = \sum_{t \in \mathbb{T}} \sum_{j \in N_t/i} u_j(y^t, x_{j0}) \quad (22)$$

The VCG transfer writes:¹⁸

$$\begin{aligned} \tau_i^{VCG}(X^t) &= \sum_{j \neq i} u_j(f^t(X^t), x_{j0}) - \sum_{j \neq i} u_j(f_{-i}^t(X^t), x_{j0}) + \\ &+ \mathbb{E}[J_{t+1}(X^{t+1}, f^t(X^t)) | X^t] - \mathbb{E}[J_{t+1}(X^{t+1}, f_{-i}^t(X^t)) | X^t], \end{aligned} \quad (23)$$

$$\forall t \in \mathbb{T}, \quad \forall i \in N_t.$$

Observe that because of independence the payoff-type component x_{i0} does not affect the distribution of X_{t+1} , thus $\mathbb{E}[J_{t+1}(X^{t+1}, y^t) | X^t]$ is invariant in x_{i0} .

We have the following result:

Lemma 1 For arbitrary period $t \in \mathbb{T}$, history X^{t-1} , agent $i \in N_t$, and payoff-type x_{i0} , the expected sum of i 's utility and the VCG transfer (23) conditional on i 's information at t is maximized if i reveals x_{i0} truthfully.

The lemma states that the optimal choice of type report by agent $i \in N_t$ is to tell the truth about his payoff type, regardless of the reported history. The proof demonstrates that given the VCG transfer the agent's problem becomes equivalent to the total expected welfare maximization. It follows that if the principal uses the accurate prior, then efficient implementation is achieved.

¹⁸Bergemann and Välimäki (2010) construct a similar VCG transfer, however in a setting with persistent population and dynamic information. In their paper, similarly to the present one, the VCG transfer generates no deficit and satisfies the participation constraint.

4.2 Scoring Transfer $\tau_i^{Sc.}$

The scoring transfer induces the truthful revelation of beliefs and is assigned upon state verification in the following period. To verify the report of agent i of cohort N_t we use the next cohort's report X_{t+1} . To construct the transfer, first calculate the probability of event X_{t+1} implied by i 's information. His information is composed of the history of reports X^{t-1} (the history of reports is public) plus the privately known type x_i . The conjugate update with respect to i 's information is:

$$\phi(\alpha|x_i, X^{t-1}) \quad (24)$$

thus the implied probability of the event X_{t+1} is

$$\Pr[X_{t+1}|x_i, X^{t-1}] = \int_{\mathcal{A}} \Pr[X_{t+1}|\alpha, x_i, X^{t-1}] \phi(\alpha|x_i, X^{t-1}) d\alpha \quad (25)$$

The principal assigns the scoring transfer equal to the natural¹⁹ logarithm of this probability:

$$\tau_i^{Sc.}(X^{t+1}) = \ln \Pr[X_{t+1}|x_i, X^{t-1}] \quad (26)$$

Note that the probability of a given state - report by the next generation, X_{t+1} , accounts not only for the distribution of true types and beliefs of the next generation, as specified by the hyperbelief-types, but also for strategic communication of N_{t+1} . The following lemma states that if the next generation is truth-telling, then the player i who faces the scoring transfer in expectation will report his belief truthfully.

Lemma 2 Fix $t \in \mathbb{T}$, $i \in N_t$ and suppose that the next-cohort's report X_{t+1} is truthful. Then the maximization of $\mathbb{E}[\tau_i^{sc.}(\hat{x}_i, X^{t+1})|x_i, X^{t-1}]$, $i \in N_t$ induces a truthful belief-type report: $\hat{x}_{ik} = x_{ik}$, $\forall k = 1, \dots, T - t$.

¹⁹The choice of logarithm base is arbitrary.

The proof relies on the assumption of conditional independence of private signals between cohorts. If the future reports are truthful, an agent's incentives to "match" the reported and the true distributions coincide. The shape of the scoring function provides these incentives. The agent will make his best bet given his information, that is, he will report the true draw from the unknown distribution. As an auxiliary result, we observe that the transfer given by (26) is information-optimal in the following sense. Suppose the designer can choose how much information about past reports to reveal to the arriving agents. It turns out that disclosing all past information increases the ex ante expectation of the scoring transfer. Consider the reduction of past information $X^{t-1} \in \mathbb{X}^{t-1} = \prod_{s=1}^t \prod_{k=0}^{T-s} \Theta_{s+k}$ as an orthogonal projection on space $\prod_{s \in S} \prod_{k \in K(s)} \Theta_s$ of types for some $S \subseteq \{1, 2, \dots, t\}$ and $K(s) \subseteq \{1, 2, \dots, T-t\}$. The reduction subsumes cases such as no disclosure about the past (S and $K(S)$ are empty) or the disclosure of only the previous generation's report. Then we have the following lemma.

Lemma 3 For all S and $K(S)$, the unconditional expectation of the reduced-information transfer is lower than the unconditional expectation of the full-information transfer:

$$\mathbb{E}\tau_i^{sc.}(x_i, proj_{S, K(S)}(X^{t-1})) \leq \mathbb{E}\tau_i^{sc.}(x_i, X^{t-1}) \quad (27)$$

Furthermore, if the projection $proj_{S, K(S)}$ is different from identity then the inequality is strict.

To obtain the result we note that both sides of Equation (27) coincide with the Shannon measure of entropy²⁰ of probability distributions $\Pr[X_{t+1}|x_i, proj_{S, K(S)}(X^{t-1})]$ and $\Pr[X_{t+1}|x_i, X^{t-1}]$, respectively. The established properties of Shannon entropy with regard to conditional distributions produce the result. (See Appendix)

²⁰with natural logarithm used as base.

4.3 Main results

The main result of this paper is that by making transfers conditional on reports of two sequential generations, the principal can reconcile dynamic efficiency with the agents' incentive to misrepresent their information. This is achieved by the delayed verification mechanism. The following proposition states the result.

Proposition 1 For any $\varepsilon > 0$ there exists $\lambda \in \mathbb{R}_+$ such that truth-telling is an ε -equilibrium of the delayed verification mechanism. Moreover, any profitable deviation that yields payoff less than ε is in the neighborhood of truth-telling.²¹

In this sense, the delayed-verification mechanism virtually implements the efficient choice rule. The appropriate transfer scaling yields truth-telling with arbitrary precision. The proof relies on Lemmas 1 and 2, and proceeds by induction starting at truth-telling at stage T . Note that N_T 's hyperbeliefs are void, and the Vickrey-Clarke-Groves transfer induce exact truthfulness of payoff types reports. The compactness of choice set \mathbb{Y} and Lipschitz-continuity of the utility functions are the required assumptions.

The delayed verification mechanism can easily be adjusted to satisfy the participation constraints and the no deficit requirement. The change is made to the scoring transfer. We balance the budget of bets by assigning to each agent i a random player j of the same cohort and let i pay j 's scoring transfer. Fix an arbitrary permutation ρ on the set N_t , such that for all $i \in N_t$, $\rho(i) \neq i$ (ρ is a derangement). The balanced scoring transfer is defined as follows:

$$\tau_i^{Sc.B} (X^{t+1}) = \tau_i^{Sc.} (X^{t+1}) - \tau_{\rho(i)}^{Sc.} (X^{t+1}) \quad (28)$$

The *balanced delayed verification mechanism* is a dynamic mechanism (f, τ) , where the allocation choice rule $f = (f_1, f_2, \dots, f_T)$ is efficient:

²¹A strategy profile is an ε -equilibrium, if there is no player and deviation that increases the player's payoff by at least ε .

$$f_t(X^t) \in \arg \max_{y_t \in Y} \left\{ \sum_{i \in N_t} u_i(y^t, x_{i0}) + \mathbb{E} [J_{t+1}(X^{t+1}, y^t) | X^t] \right\}, \quad (29)$$

and the transfer system τ is defined by the following. For all $t \in \mathbb{T}$ and $i \in N_t$:²²

$$\tau_i(X^{t+1}) = \tau_i^{VCG}(X^t) + \lambda \tau_i^{Sc.B}(X^{t+1}). \quad (30)$$

In line with the previous literature, the budget of the VCG transfer does not generate a deficit. This property is inherited by the balanced mechanism.

Proposition 2 The balanced delayed-verification mechanism satisfies the participation constraint and generates no deficit.

See proof in the Appendix.

5 Discussion

This paper shows how virtual efficiency can be achieved in a setting with sequentially arriving agents that hold independent private values for the allocation as well as private information about the future distributions of type. Both the principal and the agents know parametric class of the payoff-type distributions, however the distribution parameters are unknown. The difference in knowledge between the principal and the agents is that the latter observe informed signals about the underlying stochastic environment. Each agent receives a series of signals, drawn independently from the future type distributions. The signals reduce uncertainty about the parameter value. The principal's objective, achieved by the present mechanism, is to elicit the signals and update the hyperbelief.

The present information model allows for various degrees of initial uncertainty. The shape of the hyperdistribution reflects the principal's guesses about the parameter, as well as the quality of his information. A high entropy hyperdistribri-

²²The scoring transfer to the last cohort N_T is set to 0.

bution implies that the principal is uninformed, whereas a low entropy (in the classic case of mechanism design, degenerate) hyperdistribution corresponds to a substantial degree of certainty about the model. In this setting the agents hold additional private knowledge about the hyperdistribution, i.e., their information is strictly superior to the principal's.

We have assumed that the quality of information is homogenous across the participants and the acquisition of information is costless. As a next step, one could study the (statistically) efficient handling of information, which may be of different quality. In particular, one could allow for different betting budgets for different players, so as to provide incentives to those participants who observe more signals (and thus hold more accurate beliefs) to distinguish themselves from those with inferior information. Differentiating the betting budget across participants can be used to stimulate the information acquisition, if the acquisition is costly.

A Appendix

A.1 Summary of notation

Agent i 's information is the following:

$$x_i = \left(\begin{array}{c} x_{i0} \equiv \theta_i \\ x_{i1} \\ \vdots \\ x_{iT-t} \end{array} \right) \left. \begin{array}{l} \text{payoff-type} \\ \text{hyperbelief-type} \end{array} \right\}$$

Bold face refers to multiple variables:

$$\int_{\mathcal{A}} = \int_{\mathcal{A}_1} \int_{\mathcal{A}_2} \cdots \int_{\mathcal{A}_T} \quad ; \quad d\tilde{\alpha} = d\tilde{\alpha}_1 d\tilde{\alpha}_2 \cdots d\tilde{\alpha}_T \quad ; \quad \mathbb{X}_t = \prod_{k=0}^{T-t} \Theta_{t+k} \quad ; \quad dX_t = \prod_{i \in N_t} \prod_{k=0}^{T-t} dx_{ik}$$

$X_{t,s} = \left(x_{1_t s-t} \quad x_{2_t s-t} \quad \cdots \quad x_{i, s-t} \quad \cdots \quad x_{|N_t|_t s-t} \right)$, the signals of cohort N_t about period s , $t < s \leq T$, where 1_t denotes agent 1 in cohort N_t , 2_t agent 2 in cohort N_t etc.

To single out agent i 's payoff-type report we use the following notation:

$$(\hat{x}_{i0}, X_{-i0}^t) \equiv \left(X_1, X_2, \dots, X_{t-1}, \left(\begin{array}{cccccc} x_{1_t 0} & x_{2_t 0} & \cdots & \hat{x}_{i0} & \cdots & x_{|N_t|_t 0} \\ x_{1_t 1} & x_{2_t 1} & \cdots & x_{i1} & \cdots & x_{|N_t|_t 1} \\ x_{1_t 2} & x_{2_t 2} & \cdots & x_{i2} & \cdots & x_{|N_t|_t 2} \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ x_{1_t(T-t)} & x_{2_t(T-t)} & \cdots & x_{i(T-t)} & \cdots & x_{|N_t|_t(T-t)} \end{array} \right) \right)$$

A.2 Proof of Claim 1

Recall from Equation (4) that the efficient level of permit :

$$y_i = \sqrt{\frac{2+h+l}{\theta_i D(1+h)}}, \tag{31}$$

for all i . Then agent i 's payoff net of the scoring payment is the following:

$$U_i + \tau_i^{CG} = - \sum_j \frac{1}{\theta_j y_j} - \sum_j y_j \frac{1+h}{2+h+l} D \quad (32)$$

$$= - \sum_j \frac{1}{\theta_j} \sqrt{\frac{\theta_j D (1+h)}{2+h+l}} - \sum_j \sqrt{\frac{2+h+l}{\theta_j D (1+h)}} \frac{1+h}{2+h+l} D \quad (33)$$

$$= -2 \sum_j \sqrt{\frac{D (1+h)}{\theta_j (2+h+l)}}, \quad (34)$$

for each i , where h is the number of low, and l the number of high signals. Equation (34) implies that all agents benefit from lower h and higher l . In case of immediate verification we have the following:

$$\tau_i^{Sc.}(\hat{x}_i) = \ln(\Pr[D|\hat{x}_{i1}])^{h-i} + \ln(\Pr[0|\hat{x}_i])^{l-i} \quad (35)$$

$$= h_{-i} \ln \Pr[D|\hat{x}_i] + l_{-i} \ln \Pr[0|\hat{x}_i] \quad (36)$$

where h_{-i} (or l_{-i}) is the number of agents, excluding i , who report high (respectively, low) signal. Since $\Pr[D|D] = \Pr[0|0] = \frac{2}{3}$, and $\Pr[0|D] = \Pr[D|0] = \frac{1}{3}$ we obtain the scoring transfers as follows. If i reports a high signal (D), then:

$$\tau_i^{Sc.}(D) = h_{-i} \ln \frac{2}{3} + l \ln \frac{1}{3}. \quad (37)$$

If i reports a low signal (0).

$$\tau_i^{Sc.}(0) = h \ln \frac{1}{3} + l_{-i} \ln \frac{2}{3}. \quad (38)$$

Thus the profile of reports $(\theta_i, 0)_{i \in N}$ (truthful payoff-type report, but “no damage” irrespective of the signal) is an equilibrium. In this equilibrium, player i 's payoff equals:

$$U_i + \tau_i^{CG} + \lambda \tau_i^{Sc.} = -2 \sum_j \sqrt{\frac{D(1+h)}{\theta_j(2+h+l)}} + \lambda \left(h_{-i} \ln \frac{2}{3} + l \ln \frac{1}{3} \right) \quad (39)$$

$$= -2 \sum_j \sqrt{\frac{D(1+n)}{\theta_j(2+n)}} + \lambda \left((n-1) \ln \frac{2}{3} \right) \quad (40)$$

The equilibrium profile $(\theta_i, 0)_{i \in N}$ Pareto dominates the truthful profile $(\theta_i, x_i)_{i \in N}$. Suppose the expected transfer if the signal is high, $x_i = D$, and player i reports truthfully:

$$\mathbb{E} \left[\tau_i^{Sc.}(\hat{X}_i) | x_i = D \right] = \sum_{h_{-i}} C_{n-1}^{h_{-i}} \left(\frac{2}{3} \right)^{h_{-i}} \left(\frac{1}{3} \right)^{n-1-h_{-i}} \left(h_{-i} \ln \frac{2}{3} + (n-1-h_{-i}) \ln \frac{1}{3} \right) \quad (41)$$

$$= \sum_{h_{-i}} C_{n-1}^{h_{-i}} \left(\frac{2}{3} \right)^{h_{-i}} \left(\frac{1}{3} \right)^{n-1-h_{-i}} (h_{-i} \ln 2 - (n-1) \ln 3) \quad (42)$$

$$= (n-1) \left(\frac{2}{3} \ln 2 - \ln 3 \right) \quad (43)$$

Similarly, we can show that the expected transfer if the signal is low, $x_i = 0$ also equals $(n-1) \left(\frac{2}{3} \ln 2 - \ln 3 \right)$. Since $(n-1) \ln \frac{2}{3} > (n-1) \left(\frac{2}{3} \ln 2 - \ln 3 \right)$, player i 's utility when profile $(\theta_i, 0)_{i \in N}$ is played is greater than his utility in the truth-telling equilibrium; this holds for all i , thus the distortionary equilibrium Pareto dominates (for all $\lambda > 0$). ■

A.3 Proof of Claim 2

Consider the maximization of the expected scoring transfer by agent i . The transfer writes:

$$\int_{[0,1]} (\Pr [D|\tilde{\alpha}] \tau_i^{sc.} (\hat{x}_i, \theta_d^H) + \Pr [0|\tilde{\alpha}] \tau_i^{sc.} (\hat{x}_i, \theta_d^L)) \phi (\tilde{\alpha}|x_i) d\tilde{\alpha} \quad (44)$$

$$= \int_{[0,1]} \left(\tilde{\alpha} \ln \int_{[0,1]} \alpha \phi (\alpha|\hat{x}_i) d\alpha + (1 - \tilde{\alpha}) \ln \int_{[0,1]} (1 - \alpha) \phi (\alpha|\hat{x}_i) d\alpha \right) \phi (\tilde{\alpha}|x_i) d\tilde{\alpha} \quad (45)$$

$$= \int_{[0,1]} \tilde{\alpha} \phi (\tilde{\alpha}|x_i) \ln \int_{[0,1]} \alpha \phi (\alpha|\hat{x}_i) d\alpha + \int_{[0,1]} (1 - \tilde{\alpha}) \phi (\tilde{\alpha}|x_i) d\tilde{\alpha} \ln \int_{[0,1]} (1 - \alpha) \phi (\alpha|\hat{x}_i) d\alpha \quad (46)$$

Equation (46) boils down to:

$$p \ln \hat{p} + (1 - p) \ln (1 - \hat{p}), \quad (47)$$

where p is the probability agent i attaches to the high damage realization, \hat{p} is the probability implied by his report \hat{x}_i , and $(1 - p)$ and $(1 - \hat{p})$ are the respective complementary probabilities. The first-order condition writes $\frac{p}{\hat{p}} = \frac{1-p}{1-\hat{p}}$ and has the solution $\hat{p} = p$. The second-order condition holds. This implies that the reported signal is true at the optimum, $\hat{x}_i = x_i$. ■

A.4 Proof of Lemma 1

The problem of agent $i \in N_{t-1}$ writes:

$$\max_{\hat{x}_{i0} \in \Theta_t} \mathbb{E} \left[\underbrace{u_i (f^t (\hat{x}_{i0}, X_{-i0}^t), x_{i0}) + \tau_i^{CG} (\hat{x}_{i0}, X_{-i0}^t)}_{=: U_i (\hat{x}_{i0}, X_{-i0}^t)} \mid x_i, X^{t-1} \right], \quad (48)$$

where

$$\tau_i^{CG}(X^t) = \sum_{j \neq i} u_j(f^t(X^t), \hat{x}_{j0}) - \sum_{j \neq i} u_j(f^t(X_{-i}^t), x_{j0}) + \quad (49)$$

$$+ \mathbb{E}[J_{t+1}(X^{t+1}, f^t(X^t)) | X^t] - \mathbb{E}[J_{t+1}(X^{t+1}, f^t(X_{-i}^t)) | X^t] \quad (50)$$

eliminate the components of transfer invariant in $\hat{\theta}_i$, the problem becomes equivalent to:

$$\max_{\hat{x}_{i0} \in \Theta_i} \left\{ \mathbb{E} \left[\sum_j u_j(f^t(\hat{x}_{i0}, X_{-i0}^t), x_{j0}) + J_{t+1}(X^{t+1}, f^t(\hat{x}_{i0}, X_{-i0}^t)) | X^t \right] \right\} \quad (51)$$

Recall that for all X^t :

$$f^t(X^t) \in \arg \max_{y_t \in \mathbb{Y}} \left\{ \sum_j u_j(y_t; x_{j0}) + \mathbb{E}[J_{t+1}(X^{t+1}, y_t) | X^t] \right\} \quad (52)$$

Thus

$$x_{i0} \in \arg \max_{\hat{x}_{i0} \in \Theta_i} \left\{ \sum_j u_j(f^t(\hat{x}_{i0}, X_{-i0}^t), x_{j0}) + \mathbb{E}[J_{t+1}(X^{t+1}, f^t(\hat{x}_{i0}, X_{-i0}^t)) | X^t] \right\} \quad (53)$$

$$= \arg \max_{\hat{x}_{i0} \in \Theta_i} \mathbb{E}[U_i(\hat{x}_{i0}, X_{-i0}^t) | X^t] \quad (54)$$

for all X^t and by the law of iterated expectations (X^t contains strictly more information than x_i, X^{t-1}), hence

$$= \arg \max_{\hat{x}_{i0} \in \Theta_i} \mathbb{E}[U_i(\hat{x}_{i0}, X_{-i0}^t) | x_i, X^{t-1}] \quad (55)$$

implying that truthful report $\hat{x}_{i0} = x_{i0}$ is the solution to the initial maximization problem. The strict convexity of the utility functions yields uniqueness of the solution. \blacksquare

A.5 Proof of Lemma 2

We need to show that:

$$(x_{i1}, x_{i2}, \dots, x_{iT-t}) \in \arg \max_{(\hat{x}_{i1}, \hat{x}_{i2}, \dots, \hat{x}_{iT-t})} \mathbb{E} [\tau_i^{Sc.}(\hat{x}_i, X^{t+1}) | x_i, X^t] \quad (56)$$

where $\hat{x}_i = (x_{i0}, \hat{x}_{i1}, \hat{x}_{i2}, \dots, \hat{x}_{iT-t})$. Given that the report X_{t+1} is truthful, the expectation of the scoring transfer conditional on the agent's information is:

$$\mathbb{E} [\tau_i^{Sc.}(\hat{x}_i, X^{t+1}) | x_i, X^t] = \int_{\mathbb{X}_{t+1}} \Pr [X_{t+1} | x_i, X^{t-1}] \ln \Pr [X_{t+1} | \hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} \quad (57)$$

The proof relies on the independence of value/signal draws within and between cohorts,

$$\Pr [X_{t+1} | x_i, X^{t-1}] = \prod_{j \in N_{t+1}} \prod_{k=0,1,\dots,T-(t+1)} \Pr [x_{jk} | \hat{x}_i, X^{t-1}], \quad (58)$$

$$\Pr [x_{jk} | \hat{x}_i, X^{t-1}] = \Pr [x_{jk} | \hat{x}_{i,k+1}, X^{t-1}]. \quad (59)$$

Use Equations (58) and (59) to simplify the expected scoring transfer as follows:

$$\mathbb{E} \tau_i^{Sc.}(x_i, X^{t+1}) = \int_{\mathbb{X}_{t+1}} \Pr [X_{t+1} | x_i, X^{t-1}] \ln \Pr [X_{t+1} | \hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} \quad (60)$$

$$= \int_{\mathbb{X}_{t+1}} \Pr [X_{t+1} | x_i, X^{t-1}] \sum_{j \in N_{t+1}} \sum_{k=0,1,\dots,T-(t+1)} \ln \Pr [x_{jk} | \hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} \quad (61)$$

$$= \sum_j \sum_k \int_{\mathbb{X}_{t+1}} \Pr [X_{t+1} | x_i, X^{t-1}] \ln \Pr [x_{jk} | \hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} \quad (62)$$

$$= \sum_j \sum_k \int_{\mathbb{X}_{t+1}} \Pr [X_{t+1} | x_i, X^{t-1}] \ln \Pr [x_{jk} | \hat{x}_{i,k+1}, X^{t-1}] \mathbf{d}X_{t+1} \quad (63)$$

Then the first-order condition with respect to $\hat{x}_{i,k+1}$ writes:

$$\frac{\partial \mathbb{E} \tau_i^{Sc.}}{\partial \hat{x}_{i,k+1}} = \int_{\mathbb{X}_{t+1}} \frac{\Pr [X_{t+1}|x_i, X^{t-1}]}{\Pr [x_{jk}|\hat{x}_{i,k+1}, X^{t-1}]} \frac{\partial \Pr [x_{jk}|\hat{x}_{i,k+1}, X^{t-1}]}{\partial \hat{x}_{i,k+1}} \mathbf{d}X_{t+1} \quad (64)$$

$$= \int \frac{\Pr [x_{jk}|x_{i,k+1}, X^{t-1}]}{\Pr [x_{jk}|\hat{x}_{i,k+1}, X^{t-1}]} \frac{\partial \Pr [x_{jk}|\hat{x}_{i,k+1}, X^{t-1}]}{\partial \hat{x}_{i,k+1}} dx_{jk} = 0$$

(FOC)

The differentiation of $\int \Pr [x_{jk}|\hat{x}_{i,k+1}, X^{t-1}] dx_{jk} \equiv 1$ yields

$$\int \frac{\partial \Pr [x_{jk}|\hat{x}_{i,k+1}, X^{t-1}]}{\partial \hat{x}_{i,k+1}} dx_{jk} = 0. \quad (65)$$

Thus $\hat{x}_{i,k+1} \equiv x_{i,k+1}$, for $k = 0, 1, \dots, T - (t + 1)$, satisfy the first order condition (put differently, $\hat{x}_{i1} = x_{i1}, \hat{x}_{i2} = x_{i2}, \dots, \hat{x}_{iT-t} = x_{iT-t}$ is a critical point). To verify the second-order condition, we write (FOC) in vector form:

$$\int_{\mathbb{X}_{t+1}} \frac{\Pr [X_{t+1}|x_i, X^{t-1}]}{\Pr [X_{t+1}|\hat{x}_i, X^{t-1}]} \nabla_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} = 0 \quad (66)$$

where $\nabla_{\hat{x}_i} \Pr [\cdot]$ denotes the gradient of function $\Pr [\cdot]$ with respect to variables $\hat{x}_{i1}, \hat{x}_{i2}, \dots, \hat{x}_{iT-t}$ (column-vector), and $\mu = (\mu_{jk})$.

The second order condition of the maximization problem requires that the matrix $\nabla_{\hat{x}_i} (FOC)^T$ is negative semi-definite at the critical point.

$$\begin{aligned} \nabla_{\hat{x}_i} (FOC)^T &= \\ &= \int_{\mathbb{X}_{t+1}} -\frac{\Pr [X_{t+1}|x_i, X^{t-1}]}{(\Pr [X_{t+1}|\hat{x}_i, X^{t-1}])^2} \nabla_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}] (\nabla_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}])^T + \\ &\quad + \frac{\Pr [X_{t+1}|x_i, X^{t-1}]}{\Pr [X_{t+1}|\hat{x}_i, X^{t-1}]} H_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1}, \end{aligned} \quad (67)$$

where $H_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}]$ is the Hessian matrix. We make the following observations with respect to the components of the integral (67):

1. $-\frac{\Pr [X_{t+1}|x_i, X^{t-1}]}{(\Pr [X_{t+1}|\hat{x}_i, X^{t-1}])^2} < 0$ (note that in our model the distributions have full support, thus the inequality here is strict).

2. $\nabla_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}] (\nabla_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}])^T$ is positive definite almost surely.
3. $\int_{\mathbb{X}_{t+1}} H_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} = 0$ due to (65).
4. $\frac{\Pr [X_{t+1}|x_i, X^{t-1}]}{\Pr [X_{t+1}|\hat{x}_i, X^{t-1}]}$ is positive and bounded, thus

$$\int_{\mathbb{X}_{t+1}} \frac{\Pr [X_{t+1}|x_i, X^{t-1}]}{\Pr [X_{t+1}|\hat{x}_i, X^{t-1}]} H_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} \leq \int_{\mathbb{X}_{t+1}} c \times H_{\hat{x}_i} \Pr [X_{t+1}|\hat{x}_i, X^{t-1}] \mathbf{d}X_{t+1} = \mathbf{0}.^{23}$$

It follows that $\nabla_{\hat{x}_i} (FOC)^T$ is negative definite, and thus $\hat{x}_{i1} = x_{i1}, \hat{x}_{i2} = x_{i2}, \dots, \hat{x}_{iT-t} = x_{iT-t}$ is a strict global maximum. ■

A.6 Proof of Lemma 3

Denote the reduction in past information: $proj_{S,K(S)}(X^{t-1}) = \overline{X^{t-1}}$. This means that we consider the case when only part of past information is disclosed to the agent. Compare the transfer with reduced information:

$$\mathbb{E}_{\mathcal{T}_i^{Sc.}}(x_i, \overline{X^{t-1}}) = \mathbb{E} \left[\mathbb{E} [\ln \Pr [X_{t+1}|x_i]] \mid x_i, \overline{X^{t-1}} \right] \quad (68)$$

with the unconditional expectation of the scoring transfer with past and present information content:

$$\mathbb{E}_{\mathcal{T}_i^{Sc.}}(x_i, X^{t-1}) = \mathbb{E} \left[\mathbb{E} [\ln \Pr [X_{t+1}|x_i, X^{t-1}]] \mid x_i, X^{t-1} \right] \quad (69)$$

The right-hand side of 69 is the negative of the Shannon entropy²⁴ of random variable X_{t+1} conditional on random variables x_i and X^{t-1} that we denote $\mathcal{H}(X_{t+1}|x_i, X^{t-1})$. Similarly, the right-hand side of Equation 68 is the negative of the Shannon entropy X_{t+1} conditional on x_i only, denoted $\mathcal{H}(X_{t+1}|x_i)$. This implies that:

²³Here $\mathbf{0}$ is a matrix of zeros and respectively \leq and $=$ are componentwise (in-)equalities.

²⁴Here the entropy is defined with the natural number e as the logarithm base. Choosing a different base does not change our analysis: the constant would cancel out in Equation 70.

$$\mathbb{E}\tau_i^{Sc.}(x_i, X^{t-1}) - \mathbb{E}\tau_i^{Sc.}(x_i) = \mathcal{H}(X_{t+1}|x_i, \overline{X^{t-1}}) - \mathcal{H}(X_{t+1}|x_i, X^{t-1}) > 0 \quad (70)$$

Observe that X^{t-1} contains information about X_{t+1} that is not contained in $\overline{X^{t-1}}$ (see Equation 12), therefore the entropy of X_{t+1} conditional on X^{t-1} and x_i is lower than the entropy of X_{t+1} conditional on x_i and $\overline{X^{t-1}}$. This can also be observed from re-writing 70 in terms of unconditional entropies:

$$\begin{aligned} & \mathcal{H}(X_{t+1}|x_i, \overline{X^{t-1}}) - \mathcal{H}(X_{t+1}|x_i, X^{t-1}) = \\ & = \left(\mathcal{H}(X_{t+1}, x_i, \overline{X^{t-1}}) - \mathcal{H}(x_i, \overline{X^{t-1}}) \right) - \left(\mathcal{H}(X_{t+1}, x_i, X^{t-1}) - \mathcal{H}(x_i, X^{t-1}) \right) > 0 \end{aligned} \quad (71)$$

The marginal increase in entropy due to the addition of X_{t+1} is greater when X_{t+1} is added to x_i and $\overline{X^{t-1}}$ than when X_{t+1} is added to x_i and X^{t-1} , since the latter pair has greater informational content. Thus we obtain that:

$$\mathbb{E}\tau_i^{Sc.}(x_i, X^{t-1}) > \mathbb{E}\tau_i^{Sc.}(x_i, \overline{X^{t-1}}) \quad (72)$$

The analysis is equivalent for arbitrary projections $proj_{S, K(S)}(X^{t-1})$ different from the identity. Using the most information available increases the expected payoff. ■

A.7 Proof of Proposition 1

The proof is by induction. The inductive hypothesis for $t = 1, 2, \dots, T - 1$ is the following: If N_t report their types and beliefs truthfully, then N_{t-1} report truthfully, too.

At the last stage $t = T$ the belief-type reports are void, and the payoff-type reports are truthful due to the ex-post VCG transfer. Thus at $t = T$ there is truthful revelation (trivially for the belief-type), and therefore it suffices to prove the in-

ductive hypothesis.

Fix t . By Lemma 2 no player $i \in N_t$ can increase his *scoring* transfer by reporting anything different from the true type x_i . Thus the profit can be generated only in the remaining “welfare” part:

$$u_i(f^t(X^t), x_{i0}) + \tau_i^{CG}(X^t) =: w_i(X^t) \quad (73)$$

$w_i(X^t)$ is the total welfare change due to i 's report. Lemma 1 states that $w_i(X^t)$ is maximized at θ_i , for any belief report $\hat{x}_{i1}, \hat{x}_{i2}, \dots, \hat{x}_{iT-t}$. Thus we restrict attention to deviations in the belief report only.

If X^t is the truthful report profile, let $D_i(X^t)$ denote a transformation of X^t that replaces the hyperbelief of agent i by some $\hat{x}_{i1}, \hat{x}_{i2}, \dots, \hat{x}_{iT-t}$ different from $x_{i1}, x_{i2}, \dots, x_{iT-t}$. Denote the class of such transformations \mathcal{D}_i and let $\mathcal{D}_i^\varepsilon$ consist of all $D_i(X^t)$ that induce a welfare change larger than ε :

$$\mathcal{D}_i^\varepsilon = \{D_i \in \mathcal{D}_i : w_i(D_i(X^t)) - w_i(X^t) \geq \varepsilon\} \quad (74)$$

Consider a deviation of player i , where he distorts his hyperbelief. Under the equilibrium assumption, the profile of reports becomes $D_i(X^t) =: \hat{X}^t$. Since the choice set \mathbb{Y} is compact, the change in the allocation $f^t(\cdot)$ is bounded, implying that $\exists c \in \mathbb{R}$:

$$\|f^t(X^t) - f^t(\hat{X}^t)\| < c \quad (75)$$

By the assumption that the utility functions u_i are Lipschitz-continuous in y^t , for all i there exists K_i such that

$$\left| u_i(f^t(X^t); x_{i0}) - u_i(f^t(\hat{X}^t); x_{i0}) \right| < K_i \|f^t(X^t) - f^t(\hat{X}^t)\| \quad (76)$$

and thus $\left| u_i(f^t(X^t); x_{i0}) - u_i(f^t(\hat{X}^t); x_{i0}) \right| < cK_i$. Let $K := \max \{K_i\}_{i \in N^T/N^{t-1}}$.

The welfare writes:

$$\begin{aligned}
w_i(\hat{X}^t) &= \sum_j u_j(f^t(\hat{X}^t), x_{j0}) - \sum_{j \neq i} u_j(f_{-i}^t(\hat{X}^t), x_{j0}) + \\
&+ \mathbb{E} \left[J_{t+1}(X^{t+1}, f^t(\hat{X}^t)) \mid \hat{X}^t \right] - \mathbb{E} \left[J_{t+1}(X^{t+1}, f_{-i}^t(\hat{X}^t)) \mid \hat{X}^t \right]. \tag{77}
\end{aligned}$$

Recall that the J_t is a weighted sum of the future utilities. By a version of the Cauchy-Bunyakovsky inequality:

$$\left| \mathbb{E} \left[J_{t+1}(X^{t+1}, f^t(\hat{X}^t)) \mid \hat{X}^t \right] - \mathbb{E} \left[J_{t+1}(X^{t+1}, f^t(X^t)) \mid X^t \right] \right| \leq \tag{78}$$

$$\begin{aligned}
&\leq \int \left| \Pr[X^T \mid \hat{X}^t] - \Pr[X^T \mid X^t] \right| \times \\
&\times \sum_{j \in N^T/N^t} \left| \bar{u}_j(f^T(X^T); x_{i0}) - \bar{u}_j(f^T(\hat{X}^T); x_{i0}) \right| \mathbf{d}X^T < \tag{79}
\end{aligned}$$

$$< 2(|N^T| - |N^t|) cK, \tag{80}$$

where $\bar{u}_j(y^T; x_{i0}) = u_j(y^s; x_{i0})$ for $j \in N_s$. Similarly, for $\mathbb{E} \left[J_{t+1}(X^{t+1}, f_{-i}^t(\hat{X}^t)) \mid \hat{X}^t \right]$ we obtain

$$\begin{aligned}
&\left| \mathbb{E} \left[J_{t+1}(X^{t+1}, f_{-i}^t(\hat{X}^t)) \mid \hat{X}^t \right] - \mathbb{E} \left[J_{t+1}(X^{t+1}, f_{-i}^t(X^t)) \mid X^t \right] \right| < \\
&< 2(|N^T| - |N^t|) cK. \tag{81}
\end{aligned}$$

Summing up differences in the four components we derive that

$$\begin{aligned}
&\left| w_i(\hat{X}^t) - w_i(X^t) \right| < 4|N_t| cK + 4(|N^T| - |N^t|) cK \\
&= 4cK(|N^T| - |N^{t-1}|). \tag{82}
\end{aligned}$$

We have shown that the gain that i can achieve by misreporting the hyperbelief is bounded. Thus $\left|w_i(\hat{X}^t) - w_i(X^t)\right|$ is bounded for all $\hat{X}^t = D_i(X^t)$, $D_i \in \mathcal{D}_i^\varepsilon$. Note that, on the other hand, for any $\hat{X}^t = D_i(X^t)$, $D_i \in \mathcal{D}_i^\varepsilon$ the (negative) change in the scoring transfer, $\mathbb{E}_{\mathcal{T}_i^{Sc.}}(\hat{X}^t, X_{t+1}) - \mathbb{E}_{\mathcal{T}_i^{Sc.}}(X^{t+1})$ is also bounded from above. Thus we can choose λ such that the (negative) change in the scoring transfer, $\mathbb{E}_{\mathcal{T}_i^{Sc.}}(\hat{X}^t, X_{t+1}) - \mathbb{E}_{\mathcal{T}_i^{Sc.}}(X^{t+1})$ is less than $-4cK(|N^T| - |N^{t-1}|)/\lambda$ for any $\hat{X}^t = D_i(X^t)$, $D_i \in \mathcal{D}_i^\varepsilon$, implying that \hat{X}^t is not a profitable deviation. Going through all x_i , i and t choose the maximal λ . The maximum exists, since $|N^T|$ is finite. The inductive hypothesis is proven, hence the proposition. \blacksquare

A.8 Proof of Proposition 2

Participation constraint

The proof that the balanced mechanism satisfies the participation constraint proceeds in two steps. At the first step we show that the sum of the agent's utility and the VCG transfer is greater or equal to zero. At the second step we show that the expectation of the balanced scoring transfer equals zero, such that the entire participation constraint holds.

Step 1. By the construction of the efficient, or welfare-maximizing choice rules f and f_{-i} (see Equation (18)), it holds for all t , X^t and $i \in N_t$:

$$\begin{aligned} & \sum_j u_j(f^t(X^t), x_{j0}) + \mathbb{E}[J_{t+1}(X^{t+1}, f^t(X^t)) | X^t] \geq \\ & \sum_{j \neq i} u_j(f_{-i}^t(X^t), x_{j0}) + \mathbb{E}[J_{t+1}(X^{t+1}, f_{-i}^t(X^t)) | X^t]. \end{aligned} \quad (83)$$

Recall the definition of the Groves transfer (Equation (23) in text):

$$\begin{aligned}
\tau_i^{VCG}(X^t) &= \sum_{j \neq i} u_j(f^t(X^t), x_{j0}) - \sum_{j \neq i} u_j(f_{-i}^t(X^t), x_{j0}) + \\
&+ \mathbb{E}[J_{t+1}(X^{t+1}, f^t(X^t)) | X^t] - \mathbb{E}[J_{t+1}(X^{t+1}, f_{-i}^t(X^t)) | X^t], \tag{84}
\end{aligned}$$

Thus, the inequality (83) is equivalent to:

$$u_i(f^t(X^t), x_{i0}) + \tau_i^{VCG}(X^t) \geq 0. \tag{85}$$

By the law of iterated expectations:

$$\mathbb{E}[u_j(f^t(X^t), x_{j0}) + \tau_i^{CG}(X^t) | X^{t-1}] \geq 0. \tag{86}$$

Step 2. First, observe that:

$$\begin{aligned}
\Pr[X_{t+1}, x_j | X^{t-1}] &= \int_{\mathcal{A}} \Pr[X_{t+1}, x_j | \alpha, X^{t-1}] \phi(\alpha | X^{t-1}) d\alpha = \\
&= \int_{\mathcal{A}} \Pr[X_{t+1} | \alpha, X^{t-1}] \Pr[x_j | \alpha, X^{t-1}] \phi(\alpha | X^{t-1}) d\alpha. \tag{87}
\end{aligned}$$

$$\Pr[x_j | X^{t-1}] = \int_{\mathcal{A}} \Pr[x_j | \alpha, X^{t-1}] \phi(\alpha | X^{t-1}) d\alpha. \tag{88}$$

Since the signals and types are drawn independently (from the true α -distributions), $\Pr(x_j | \alpha, X^{t-1})$ both expressions (87) and (88) are invariant in j . Therefore

$$\begin{aligned}
&\ln \Pr[X_{t+1}, x_j | X^{t-1}] - \ln \Pr[x_j | X^{t-1}] \\
&= \ln \Pr[X_{t+1} | x_j, X^{t-1}] = \tau_j^{Sc.}(X^{t+1}) \tag{89}
\end{aligned}$$

is also invariant in j . This implies that the expectation of the own and the opponent's transfer are equal, conditional on the past history X^{t-1} :

$$\mathbb{E} [\tau_i^{Sc.} (X^{t+1}) | X^{t-1}] = \mathbb{E} [\tau_{\rho(i)}^{Sc.} (X^{t+1}) | X^{t-1}]. \quad (90)$$

Thus the expectation of the balanced scoring transfer is zero:

$$\mathbb{E} [\tau_i^{Sc.B} (X^{t+1}) | X^{t-1}] = \mathbb{E} [\tau_i^{Sc.} (X^{t+1}) - \tau_{\rho(i)}^{Sc.} (X^{t+1}) | X^{t-1}] = 0. \quad (91)$$

We conclude that

$$\mathbb{E} [u_i (f^t (X^t), x_{j0}) + \tau_i^{CG} (X^t) + \tau_i^{Sc.B} (X^{t+1}) | X^{t-1}] \geq 0, \quad (92)$$

thus the participation constraint holds. ■(PC)

No deficit

From Equation (83) it immediately follows that the Vickrey-Clarke-Groves transfer to each agent is non-positive:

$$\tau_i^{CG} (X^t) \leq 0$$

The balanced scoring transfer satisfies $\sum_{i \in N_t} \tau_i^{Sc.B} (X^{t+1}) = 0$ by construction (See Equation (28)). Therefore, the aggregate no-deficit constraint holds:

$$\sum_{i \in N_t} \tau_i^{CG} (X^t) + \tau_i^{Sc.B} (X^{t+1}) \leq 0. \quad \blacksquare(BB)$$

References

- Athey, S., Segal, I., 2013. An efficient dynamic mechanism. *Econometrica* 81 (6), 2463–2485. 2, 3, 11
- Bellman, R., 1966. *Dynamic Programming*. *Science* 153 (3731), 34–37.
 URL <http://www.sciencemag.org/cgi/content/abstract/153/3731/34> 14

- Bergemann, D., Välimäki, J., 2010. The dynamic pivot mechanism. *Econometrica* 78 (2), 771–789. 2, 3, 18
- Cavallo, R., Parkes, D. C., Singh, S., 2009. Efficient mechanisms with dynamic populations and dynamic types. Tech. rep., Harvard University. Division of Engineering and Applied Physics. 2
- Clarke, E., 1971. Multipart pricing of public goods. *Public choice* 11 (1), 17–33. 3
- Crémer, J., McLean, R. P., 1985. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* 53 (2), pp. 345–361.
URL <http://www.jstor.org/stable/1911240> 4, 5, 8
- Crémer, J., McLean, R. P., 1988. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica* 56 (6), pp. 1247–1257.
URL <http://www.jstor.org/stable/1913096> 4
- d’Aspremont, C., Crémer, J., Gérard-Varet, L.-A., 2004. Balanced bayesian mechanisms. *Journal of Economic Theory* 115 (2), 385 – 396.
URL <http://www.sciencedirect.com/science/article/B6WJ3-49RCDWB-1/2/5564cdf55764d454291cbf824f790c63> 3
- Friedman, E. J., Parkes, D. C., 2003. Pricing wifi at starbucks: issues in online mechanism design. In: *Proceedings of the 4th ACM conference on Electronic commerce*. ACM, pp. 240–241. 3
- Gershkov, A., Moldovanu, B., 2009. Learning about the future and dynamic efficiency. *The American Economic Review*, 1576–1587. 2
- Groves, T., 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* 41 (4), 617–631. 3
- Hanson, R., 2013. Shall we vote on values, but bet on beliefs? *Journal of Political Philosophy*, n/a–n/a.
URL <http://dx.doi.org/10.1111/jopp.12008> 17

- Lavi, R., Nisan, N., 2004. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science* 310 (1 - 3), 159 – 180.
URL <http://www.sciencedirect.com/science/article/pii/S0304397503003918> 2
- McAfee, R., Reny, P., 1992. Correlated information and mechanism design. *Econometrica: Journal of the Econometric Society*, 395–421. 4
- Mezzetti, C., 2004. Mechanism design with interdependent valuations: Efficiency. *Econometrica* 72 (5), 1617–1626. 2
- Parkes, D., Singh, S., 2003. An MDP-based approach to online mechanism design. In: *Proc. 17th Annual Conf. on Neural Information Processing Systems (NIPS 2003)*. Citeseer. 2
- Pavan, A., Segal, I., Toikka, J., 2010. Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure. *Profit Maximization and Information Disclosure*. 3
- Raiffa, H., Schleifer, R., 1961. *Applied statistical decision theory*. 2, 3
- Vickrey, W., 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16 (1), 8–37. 3