Stylized Implementation

John Y. Zhu[1]

November 10, 2021

Abstract

In many implementation problems, one does not know the true probability governing the type space Ω but will know some summary statistics. To this setting I introduce *stylized implementation*: The mechanism designer first whittles Ω down to a high probability event Ω^*. She then ex-post implements the desired decision over Ω^*, yielding some mechanism as the solution. I argue that if the mechanism designer uses that mechanism, then, with high probability, the desired decision will be implemented over Ω. Taking the stylized approach to implementation can yield significantly better solutions than taking the ex-post approach. An application to repeated resource allocation is considered.

JEL Codes:
Keywords:

[1]University of Kansas, johnzhuyiran@ku.edu.

1 Introduction

In a typical implementation problem involving a mechanism designer and a set of agents, it is unlikely that the true probability measure governing the type space Ω will be common knowledge, or even known to anyone. On the other hand, having common knowledge of some basic summary statistics about Ω seems quite plausible. Such knowledge could come from prior experience dealing with similar decision problems or groups of players.

In such a scenario, without a prior, the mechanism designer is unable to take the Bayesian approach to implementation. Ex-post implementation is certainly an option, but that would involve completely ignoring the available summary statistics. Intuitively, ignoring such statistics could be quite suboptimal. For example, if a type involves repeated draws from a distribution – a common situation in dynamic decision problems – and it is known that the draws are close to independent, then when there are many dates, the Central Limit Theorem would imply that the realized type is highly likely to lie in a tiny sliver of the type space. Such a seemingly important fact would simply not be incorporated in ex-post implementation.

In this paper, I present an alternative approach to implementation, related to ex-post implementation, but that can incorporate those potentially very valuable summary statistics.

To begin with, I introduce a new equilibrium notion weaker than ex-post equilibrium: Given a direct mechanism, possibly dynamic, I define what it means for truth-telling to be an ε-ex-post equilibrium. When truth-telling is an ε-ex-post equilibrium, I argue that one can expect, with high probability, all agents to report the truth at all dates.

And now, consider the following *stylized approach to implementation*: The mechanism designer whittles the true type space Ω down to a smaller "stylized type space" Ω^*. The whittling must be guided by the known summary statistics: Ω^* needs to be chosen in a way so that, given the known summary statistics, it is common knowledge that the realized type ω will land in Ω^* with sufficiently high probability.

The mechanism designer then takes the ex-post approach to solve the implementation problem over Ω^*. Let M^*, a direct mechanism over Ω^*, be a resulting solution. Compose M^* with a retraction mapping $[\cdot] : \Omega \to \Omega^*$ to get a direct mechanism $M^* \circ [\cdot]$ over the true type space Ω. Call it a stylized mechanism.

The stylized approach to implementation is to generate such a stylized mechanism and use it on the agents.

In what sense does this approach "work?" In the paper I show that truth-telling is an ε-ex-post equilibrium of $M^* \circ [\cdot]$. Consequently, the mechanism designer can expect with high probability that all agents report the truth at all dates. Moreover, recall, M^* implements the desired decision over Ω^*, which is, itself, a high probability event. Putting these two together, and we conclude that $M^* \circ [\cdot]$ implements the desired decision with high probability.

If the mechanism designer is comfortable with implementation with high probability rather than implementation over the entire type space, then taking the stylized approach to implementation can yield significantly better mechanisms than taking the ex-post approach. In particular, in settings where there is a notion of cost, it can yield significantly cheaper mechanisms.

In the second half of the paper I demonstrate this by considering an application to a repeated resource allocation problem. The setting is quasilinear, agents are protected by limited liability, and the mechanism designer can make nonnegative transfers to the agents in an effort to implement the efficient allocation of resources each date. I show that when the number of agents and dates goes to infinity, the cheapest ex-post mechanism – which is essentially just a sequence of VCG mechanisms – has an infinite cost-to-surplus ratio. On the other hand, if agents are patient and some "Central Limit Theorem style" summary statistics are known, then, by taking the stylized approach to implementation, the principal can – via a stylized mechanism I call the *linked VCG mechanism* – implement the efficient allocation almost surely at a cost-to-surplus ratio of zero.

Recently, Lee (2017) and Azevedo and Budish (2019) have explored notions of approximate strategy-proofness. My work is related in the sense that a stylized mechanism can be viewed as being an approximately ex-post mechanism. Also related are Jackson and Manelli (1997), who show how the attractive properties of the market mechanism are approximately preserved when the price-taking assumption is relaxed, and Bergemann and Välimäki (2002), who consider dynamic ex-post implementation when there is a common prior. The concept of ε-ex-post equilibrium is related to the contemporaneous perfect ε-equilibrium of Mailath, Postlewaite, and Samuelson (2005).

In the application to repeated resource allocation, my work on linking VCG mechanisms over stylized type spaces is related the work of Holmström (1979) on VCG mechanisms over restricted preference domains. See also Green and Laffont (1977). In the many agents and dates limit, the linked VCG mechanism that implements the efficient allocation almost surely at a cost-to-surplus ratio of zero can be viewed as a type of budget mechanism. A number of papers have shown how budget mechanisms can align incentives across multiple problems when transfers are unavailable. See, for example, Jackson and Sonnenschein (2007) and Frankel (2014). My work reveals a surprising connection between budget mechanisms and VCG mechanisms.

2 Stylized Implementation

Decision Problems.

An N-agent T-date *decision problem* is a triple (Ω, D, U). $\Omega = \Pi_{1\leq n\leq N,\ 1\leq t\leq T}\Omega_t^n$ is the type space, where each Ω_t^n is a finite set of date t types for agent n. $D = \Pi_{t=1}^T D_t$ is a finite set of decision sequences. $U^n : D \times \Omega \to \mathbb{R}$ is agent n's payoff function –

which can depend on other agents' types – and is defined to be

$$U^n(d,\omega) = \sum_{t=1}^{T} \beta^{t-1} u_t^n(d|_t, \omega|_t),$$

where $\beta \in (0,1]$ is the discount factor and $u_t^n : D|_t \times \Omega|_t \to \mathbb{R}$ is agent n's date t utility function, which depends on the history of decisions, $d|_t$, and types, $\omega|_t$, up through date t.

For the rest of this section, fix a decision problem (Ω, D, U).

Notation.

For an object $\cdot_t^n$ indexed by agents and dates, let the superscript denote the agent index and the subscript denote the date index. Let $\cdot^n$ denote agent n's *sequence* of $\cdot_t^n$ across all dates and let $\cdot_t$ denote the date t *profile* of $\cdot_t^n$ across all agents. Let $\cdot$ denote the array of $\cdot_t^n$ across agents and dates. If an object $\cdot^n$ is only indexed by agents, then let $\cdot$ denote the profile of $\cdot^n$ across agents. If an object $\cdot_t$ is only indexed by date, then let $\cdot$ denote the sequence of $\cdot_t$ across all dates, and let $\cdot|_t$ denote the subsequence of $\cdot$ up through date t.

Let A and B be two sets of sequences. A map $f : A \to B$ is adapted is $a|_t = a'|_t \Rightarrow f(a)|_t = f(a')|_t$.

Statistical Models

A *statistical model* is summarized by a nonempty set of probabilities, $\mathcal{P}$, over Ω. It is common knowledge that ω is governed by some true probability, call it P, lying in $\mathcal{P}$. In a typical application, $\mathcal{P}$ will be an infinite set of probabilities derived from some exogenous commonly known summary statistics about ω.

In addition, I assume it is common knowledge each agent n knows his own marginal, P^n, of the true probability.

Direct Mechanisms

A direct mechanism is an adapted map $M : \Omega \to D$.

Given a direct mechanism M, an agent n strategy, σ^n, consists of a sequence of maps $\sigma_t^n : D|_{t-1} \times \Omega^n|_t \to \Omega_t^n$. Let Σ^n denote the set of all agent n strategies. A profile of strategies, σ, can be viewed as an adapted map $\sigma : \Omega \to \Omega$. Let *id* be the strategy profile in which all agents report the truth at all dates.

ε-Ex-Post Equilibrium

For the rest of the paper, fix an $\varepsilon > 0$, to be interpreted as "small."

Given a direct mechanism M, consider the following scenario: Agent n is standing at date t, having

- observed $(d|_{t-1}, \omega^n|_t)$,
- reported the truth up through date $t-1$, and
- conjectured that all other agents are playing id^{-n}.

He then wonders – if my type ends up being some ω^n, what is the maximum regret I will feel, in today's terms, if I continue to report the truth until the end?[2]

The answer is the following quantity:

$$R_t^n(d|_{t-1}, \omega^n) = \max_{\substack{\omega^{-n} \in \Omega^{-n} \ s.t.\ M(\omega^{-n}, \omega^n)|_{t-1} = d|_{t-1}, \\ \hat{\omega}^n \in \Omega^n \ s.t.\ \hat{\omega}^n|_t = \omega^n|_t}} \sum_{s=t}^{T} \beta^{s-t} [u_s^n(M(\omega^{-n}, \hat{\omega}^n)|_s, (\omega^{-n}, \omega^n)|_s) - u_s^n(M(\omega^{-n}, \omega^n)|_s, (\omega^{-n}, \omega^n)|_s)].$$

Of course, at date t, agent n has only observed $\omega^n|_t$, and so does not know what ω^n will be, and therefore what his maximum regret $R_t^n(d|_{t-1}, \omega^n)$ is. However, he does know P^n, and can therefore form a conditional expected maximum regret, $\mathbf{E}_{P^n}[R_t^n(d|_{t-1}, \omega^n) \mid \omega^n|_t]$. This quantity measures how far off from being ex-post optimal it is for agent n to report the truth from date t until the end, assuming all other agents play id^{-n}, and agent n has reported the truth up through date $t-1$.

Definition. *An agent n strategy σ^n is reasonable given the conjecture that all other agents play id^{-n} if, for all dates t,*

$$\mathbf{E}_{P^n}[R_s^n(d|_{s-1}, \omega^n) \mid \omega^n|_s] \le \varepsilon \quad \forall s \le t \Rightarrow \sigma_t^n(d|_{t-1}, \omega^n|_t) = \omega_t^n. \tag{1}$$

Consider agent n deciding what to report at date 1. Given his conjecture that all other agents play id^{-n}, if the left side of (1) is satisfied for $t = 1$, then it could be said that, from the perspective of agent n at date 1, reporting the truth starting from today is within ε of being ex-post optimal. Therefore, I assume agent n reports the truth at date 1. The definition of reasonable σ^n is then justified by induction.

Let $\Sigma^n(id^{-n})$ denote the set of all reasonable σ^n, and let $\Sigma(id)$ denote the set of all reasonable strategy profiles.

Definition. *id is an ε-ex-post equilibrium if it is common knowledge that every reasonable strategy profile, σ, satisfies $P(\sigma(\omega) \neq \omega) \le \varepsilon$.*

[2] What if the observed $(d|_{t-1}, \omega^n|_t)$ is incompatible with the conjecture that all other agents play id^{-n}, assuming agent n has reported the truth up through date $t-1$? This is impossible. If this were the case, then $(d|_{t-1}, \omega^n|_t)$ would be incompatible with any σ^{-n}, and, therefore, cannot be observed by an agent n who has reported the truth up through date $t-1$.

When id is an ε-ex-post equilibrium of a direct mechanism, I will interpret it to mean that one can expect some reasonable strategy profile $\sigma \in \Sigma(id)$ to be played. In a typical application, the set, $\Sigma(id)$, will not be known to anyone. Nevertheless, one will still be able to deduce things just from the knowledge that the strategy profile being played belongs in $\Sigma(id)$ – for example, one knows that all agents will report the truth at all dates will probability at least $1 - \varepsilon$.

I am now ready to introduce the stylized approach to implementation.

Stylized Type Spaces

For the rest of the paper, assume private values: $U^n(d, \omega) = U^n(d, \omega^n)$ for all n.

Given a nonempty subset $\Omega^{*n} \subset \Omega^n$ for each agent n, define the stylized type space $\Omega^* = \Pi_{1 \leq n \leq N} \Omega^{*n}$. A *retraction* $[\cdot]^n$ is an adapted map from Ω^n to Ω^{*n} that is the identity on Ω^{*n}.

Lemma 1. *There exists a retraction from Ω^n to Ω^{*n}.*

Proof of Lemma 1. Let $t(\omega^n)$ be the first date t for which there does not exist an $\hat{\omega}^n \in \Omega^{*n}$ such that $\hat{\omega}^n|_t = \omega^n|_t$. $t(\omega^n)$ is a stopping time. If $\omega^n \in \Omega^{*n}$, then set $t(\omega^n) = T+1$. For each $\omega^n|_{t(\omega^n)-1}$, select a $\hat{\omega}^n \in \Omega^{*n}$ such that $\hat{\omega}^n|_{t(\omega^n)-1} = \omega^n|_{t(\omega^n)-1}$. Define $[\omega^n]$ to be the $\hat{\omega}^n$ selected given $\omega^n|_{t(\omega^n)-1}$. It is clearly the identity function over Ω^{*n}.

To verify $[\cdot]^n$ is adapted, let $\omega'^n, \omega''^n \in \Omega^n$ satisfy $\omega'^n|_t = \omega''^n|_t$ for some t. Since $t(\omega)$ is a stopping time, it must be that either $t \geq t(\omega'^n) = t(\omega''^n)$ or $t < t(\omega'^n), t(\omega''^n)$. In the former case, $[\omega'^n] = [\omega''^n]$. In the latter case, $[\omega'^n]|_t = \omega'^n|_t = \omega''|_t = [\omega''^n]|_t$. □

Suppose an ex-post direct mechanism over Ω^* has been identified – that is, an adapted map $M^* : \Omega^* \to D$ satisfying

$$U^n(M^*(\omega^{-n}, \omega^n), \omega^n) \geq U^n(M^*(\omega^{-n}, \hat{\omega}^n), \omega^n) \quad \forall n,\ \omega^{-n} \in \Omega^{*-n},\ \omega^n, \hat{\omega}^n \in \Omega^{*n}.$$

Then composing M^* with a retraction profile yields a direct mechanism, $M^* \circ [\cdot]$. I will call such a direct mechanism a *stylized mechanism*, and, when denoting a stylized mechanism, I will specify both the direct mechanism and the underlying stylized type space.

I now show, given a stylized mechanism $(\Omega^*, M^* \circ [\cdot])$ and a statistical model $\mathcal{P}$, if it is common knowledge $P^n(\Omega^{*n})$ is sufficiently close to 1 for all agents n, then id is an ε-ex-post equilibrium of $M^* \circ [\cdot]$.

Stylized Mechanisms

Fix a statistical model $\mathcal{P}$, a stylized mechanism $(\Omega^*, M^* \circ [\cdot])$, and let $\overline{R}$ be an upper bound on $R^n_t(d|_{t-1}, \omega^n)$ for all n, t.

Lemma 2. *If $\omega^n \in \Omega^{*n}$ then $R_t^n(d|_{t-1}, \omega^n) = 0$.*

Proof. The proof relies on the private values assumption and the fact that the way Ω^n is retracted onto Ω^{*n} is independent of Ω^{-n} for all agents n.

Let $\omega^n \in \Omega^{*n}$ and $R_t^n(d|_{t-1}, \omega^n)$ be defined. Then

$$\begin{aligned} R_t^n(d|_{t-1}, \omega^n) &= \max_{\substack{\omega^{-n} \in \Omega^{-n} \ s.t. \ M(\omega^{-n}, \omega^n)|_{t-1} = d|_{t-1}, \\ \hat{\omega}^n \in \Omega^n \ s.t. \ \hat{\omega}^n|_t = \omega^n|_t}} \sum_{s=t}^{T} \beta^{s-t} [u_s^n(M^* \circ [\omega^{-n}, \hat{\omega}^n]|_s, \omega^n|_s) \\ &\qquad - u_s^n(M^* \circ [\omega^{-n}, \omega^n]|_s, \omega^n|_s)] \\ &= \max_{\substack{\omega^{-n} \in \Omega^{-n} \ s.t. \ M(\omega^{-n}, \omega^n)|_{t-1} = d|_{t-1}, \\ \hat{\omega}^n \in \Omega^n \ s.t. \ \hat{\omega}^n|_t = \omega^n|_t}} \beta^{-t} \left[U^n(M^*([\omega^{-n}]^{-n}, [\hat{\omega}^n]^n), \omega^n) \right. \\ &\qquad \left. - U^n(M^*([\omega^{-n}]^{-n}, \omega^n), \omega^n) \right]. \end{aligned}$$

Since $[\omega^{-n}]^{-n} \in \Omega^{*-n}$, $[\hat{\omega}^n]^n, \omega^n \in \Omega^{*n}$, and M^* is an ex-post direct mechanism over Ω^*, therefore, $\beta^{-t} [U^n(M^*([\omega^{-n}]^{-n}, [\hat{\omega}^n]^n), \omega^n) - U^n(M^*([\omega^{-n}]^{-n}, \omega^n), \omega^n)] = 0$. Thus, $R_t^n(d|_{t-1}, \omega^n) = 0$. □

Proposition 1. *For any $c > 0$, if it is common knowledge that $P^n(\omega^n \notin \Omega^{*n}) \leq \frac{c\varepsilon}{N\overline{R}}$ for all agents n, then it is common knowledge that $P(\sigma(\omega) \neq \omega) < c$ for all $\sigma \in \Sigma(id)$.*

In particular, if $c \leq \varepsilon$, then id is an ε-ex-post equilibrium.

Proof of Proposition 1. Define $X_t^n(\omega) := P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_t)$. Extend the sequence by one date by defining $X_{T+1}^n = X_T^n$. It is common knowledge X^n is a nonnegative martingale with respect to P with expected value $X_0^n = P^n(\omega^n \notin \Omega^{*n}) \leq \frac{c\varepsilon}{N\overline{R}}$.

Let τ^n denote the stopping time when X_t^n first exceeds $\frac{\varepsilon}{\overline{R}}$. If X_t^n never exceeds $\frac{\varepsilon}{\overline{R}}$, then set $\tau^n = T + 1$. Let $E^n \subset \Omega$ denote the event on which there exists a date t such that X_t^n exceeds $\frac{\varepsilon}{\overline{R}}$. By Doob's optional stopping theorem, we have

$$\frac{c\varepsilon}{N\overline{R}} \geq X_0^n = \mathbf{E}X_\tau^n = \mathbf{E}X_\tau^n 1_{\tau \leq T} + \mathbf{E}X_\tau^n 1_{\tau = T+1} \geq \mathbf{E}X_\tau^n 1_{\tau \leq T} > \frac{\varepsilon}{\overline{R}} \mathbf{E}1_{\tau \leq T} = \frac{\varepsilon}{\overline{R}} P(E^n).$$

Thus, it is common knowledge $P(E^n) < \frac{c}{N}$. Now, by de Morgan's Law, it is common knowledge that

$$P\left(\exists t, n \ P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_t) > \frac{\varepsilon}{\overline{R}}\right) = P(\cup_{n=1}^N E^n) \leq \sum_{n=1}^{N} P(E^n) < c.$$

Let $\omega \notin \cup_{n=1}^N E^n$. Then $\forall t, n \ P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_t) \leq \frac{\varepsilon}{\overline{R}}$. Since Lemma 2 implies $\mathbf{E}_{P^n} [R_t^n(d|_{t-1}, \omega^n) \mid \omega^n|_t] \leq P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_t)\overline{R}$, we have

$$\forall t, n \ \mathbf{E}_{P^n} [R_t^n(d|_{t-1}, \omega^n) \mid \omega^n|_t] \leq \varepsilon.$$

Let $\sigma \in \Sigma(id)$. Then, by definition, $\sigma(\omega) = \omega$. Thus, $P(\sigma(\omega) \neq \omega) \leq P(\cup_{n=1}^{N} E^n) < c$. $\square$

Stylized Implementation

Consider the following set of primitives:

> A decision problem (Ω, D, U), a statistical model $\mathcal{P}$, and an implementation problem consisting of, for each $\omega \in \Omega$, a set of desirable decisions $D(\omega) \subset D$.
>
> Example: In an auction decision problem, a decision would be an allocation of the object along with payments from the bidders. And now, if the implementation problem is the efficient one, then the set of desirable decisions would be those that allocate the object to the bidder with the highest valuation.

The ex-post approach to implementation is to find an ex-post direct mechanism $M : \Omega \to D$ satisfying $M(\omega) \in D(\omega)$ for all $\omega \in \Omega$.

If the mechanism designer is willing to accept implementation with high probability over implementation for all types, then the theory we have developed in this paper – in particular, Proposition 1 – provides the mechanism designer with an alternative: The stylized approach to implementation.

- First choose a stylized type space Ω^* such that it is common knowledge $P^n(\Omega^{*n})$ is sufficiently close to 1 for all agents n.
- Take the ex-post approach to solving the implementation problem over Ω^*.
- This yields some ex-post direct mechanism $M^* : \Omega^* \to D$ satisfying $M^*(\omega) \in D(\omega)$ for all $\omega \in \Omega^*$.
- Create the stylized mechanism $M^* \circ [\cdot]$. With high probability, $M^* \circ [\omega] \in D(\omega)$.

Moreover, Proposition 1 tells us the mechanism designer can control how high is the probability the desired decision is implemented by choosing how likely is Ω^{*n} for each n. In particular, as all Ω^{*n} become almost sure, $M^* \circ [\cdot]$ implements the desired decision almost surely.

3 Application: Repeated Resource Allocation

A Model of Repeated Resource Allocation.

A principal (she) possesses a quantity $\overline{q}$ of a divisible, durable resource. She repeatedly allocates this resource to a set of $N \geq 2$ agents across $T \geq 1$ dates.

At each date t, each agent n is endowed with $\omega_t^n \in [0, \infty)$ units of a project type, f. f is a strictly concave, C^1 function $f : [0, \infty) \to [0, \infty)$ that maps resource quantity to payoff. Assume $f'(0) < \infty$.

An allocation array a assigns agent n at date t an amount $a_t^n \geq 0$ of the resource, subject to feasibility constraints, $\sum_{n=1}^N a_t^n \leq \overline{q}$ for all t. A transfer profile w specifies a profile of nonnegative payments from the principal to the agents at date T.

The principal desires to efficiently allocate her resource each date.

One application of this model is to an organization's problem of designing an *internal talent marketplace*. Instead of having a static collection of employee-job matchings, many organizations are reimagining work as a flow of discrete tasks that need to be assigned to available employees through some dynamic mechanism. See Smet, Lund and Schaninger (2016).

This problem can be viewed through the repeated resource allocation model: The principal corresponds to the organization's headquarters and the agents correspond to various departments. Projects are departmental tasks. The stock of durable resources is the organization's pool of employees parameterized by hours of labor per date, where a date could be, say, one month. Transfers from the principal to agents correspond to incentive pay for department managers.

Of course, realistically, departments might have different types of projects and employees might have different skill sets that may make them more suitable to some projects than others. The repeated resource allocation model can be generalized to accommodate such realism by allowing for multiples project and resource types. The analysis below can be adapted to handle such a generalization of the model.

The Induced Decision Problem

The repeated resource allocation model defines an N-agent T-date decision problem:

- $\Omega = [0, \infty)^{NT}$,
- $D = \{(a, w) \mid \sum_{n=1}^N a_t^n \leq \overline{q} \;\; \forall t \;\text{ and }\; w^n \geq 0 \;\; \forall n\}$, and
- $U^n((a, w), \omega) = U^n((a^n, w^n), \omega^n) = \sum_{t=1}^T \beta^{t-1} \omega_t^n f\left(\frac{a_t^n}{\omega_t^n}\right) + \beta^{T-1} w^n$ for all n.

In addition, define the following auxiliary quantities,

- agent n surplus: $S^n((a, w), \omega) = S^n(a^n, \omega^n) = \sum_{t=1}^T \beta^{t-1} \omega_t^n f\left(\frac{a_t^n}{\omega_t^n}\right)$,
- total surplus: $S((a, w), \omega) = S(a, \omega) = \sum_{n=1}^N S^n(a^n, \omega^n)$, and
- cost: $C((a, w), \omega) = C(w) = \sum w^n$.

A direct mechanism can be expressed as a pair of adapted maps $(A, W) : \Omega \to D$ consisting of an allocation map and a transfer map. The efficient allocation map is the unique allocation map, $\boldsymbol{A}$, satisfying

$$\boldsymbol{A}_t^n(\omega) = \frac{\omega_t^n}{\sum_{m=1}^N \omega_t^m} \cdot \overline{v} \qquad \forall \omega \in \Omega.$$

A direct mechanism (A, W) is efficient if $A \equiv \boldsymbol{A}$.

The principal's desire to efficiently allocate her resource each date induces an implementation problem where the set of desirable decisions is $D(\omega) = \{(\boldsymbol{A}(\omega), w) \mid w^n \geq 0 \ \forall n\}$ for all $\omega \in \Omega$.

I now compare the ex-post and stylized approaches to implementation, with a focus on which approach is cheaper for the principal.

The Unlinked VCG Mechanism

Suppose the principal wants to ex-post implement the efficient allocation. One option is to run a separate Vickrey-Clark-Groves (VCG) mechanism each date, paying each agent the sum of all other agents' contributions to surplus:

Definition. *The unlinked VCG mechanism* $(\boldsymbol{A}, V)$ *is the efficient direct mechanism with transfer map defined as follows: For all* $\omega \in \Omega$*,*

$$V^n(\omega) = \sum_{m \neq n} \sum_{t=1}^{T} \beta^{t-T} \omega_t^m f\left(\frac{\boldsymbol{A}_t^m(\omega)}{\omega_t^m}\right).$$

Proposition 2. *The unlinked VCG mechanism is the cheapest efficient ex-post direct mechanism: Let* $(\boldsymbol{A}, W)$ *be any efficient ex-post direct mechanism. Then for every* $\omega \in \Omega$*, we have* $C(V(\omega)) \leq C(W(\omega))$*.*

Proposition 2 is a corollary of Proposition 3 below.

Even though the unlinked VCG mechanism is the cheapest mechanism that ex-post implements the efficient allocation, it is still expensive with cost-to-surplus ratio

$$\frac{C(\omega)}{S(\boldsymbol{A}(\omega), \omega)} = N - 1.$$

In particular, as the number of agents tends to infinity, so does the cost-to-surplus ratio.

The Linked VCG Mechanism

Definition. *Given Ω^*, the linked VCG mechanism $(\boldsymbol{A}|_{\Omega^*}, V^*)$ is an efficient ex-post direct mechanism over Ω^* with transfer map defined as follows: For all $\omega \in \Omega^*$,*

$$V^{*n}(\omega) = V^n(\omega) - \underset{\hat{\omega}^n \in \Omega^{*n}}{\arg\min} V^n(\omega^{-n}, \hat{\omega}^n).$$

Proposition 3. *If Ω^* is convex, then the linked VCG mechanism $(\boldsymbol{A}|_{\Omega^*}, V^*)$ is the cheapest efficient ex-post direct mechanism over Ω^*.*

Proposition 3 is a consequence of Theorem 1 of Hölmström (1979) about the necessity of VCG mechanisms over restricted domains. The proof is a straightforward application of the envelope theorem.

Since the decision problem is one of private values, given an arbitrary retraction profile $[\cdot] : \Omega \to \Omega^*$ (which exists by Lemma 1), Proposition 1 implies that $(\Omega^*, (\boldsymbol{A}|_{\Omega^*}, V^*) \circ [\cdot])$ implements the efficient allocation with high probability provided it is common knowledge $P^n(\Omega^{*n})$ is sufficiently close to 1 for each agent n. As an abuse of nomenclature, call $(\Omega^*, (\boldsymbol{A}|_{\Omega^*}, V^*) \circ [\cdot])$ a linked VCG mechanism as well, and from now on I will denote it by (Ω^*, V^*).

I now show, as the number of agents and dates goes to infinity, assuming agents are patient and $\mathcal{P}$ implies common knowledge of some "Central Limit Theorem style" summary statistics, then, by taking the stylized approach to implementation, the principal can implement the efficient allocation almost surely at a cost-to-surplus ratio of zero. This is in stark contrast to taking the ex-post approach, which would entail a cost-to-surplus ratio of infinity.

A Family of Repeated Resource Allocation Models

Fix a quantity $q > 0$ of the resource and a project type f. Consider the family of decision problems parameterized by N satisfying $(\overline{q}(N), f(N)) = (Nq, f)$ and $T(N) = N$. Refer to the member of the family with N-agents as the N-agent decision problem. Throughout the analysis below, we may append (N) to a parameter to emphasize that it is belongs to the N-agent decision problem.

Assume the family of statistical models $\{\mathcal{P}(N)\}_{N \geq 2}$ satisfies the following "Central Limit Theorem style" summary statistics:

Assumption 1. *There exist $\omega^{max} > \omega^{avg} > 0$ and an increasing function $I : (0, \infty) \to (0, \infty)$ satisfying $\lim_{x \to \infty} I(x) = \infty$ such that, for each N-agent decision problem, it* *is common knowledge that*

$$P(N)(\omega_t^n > \omega^{max}) = 0 \qquad \forall n, t \leq N,$$

$$P(N)\left[\left|\frac{\sum_{t=1}^N \omega_t^n}{N} - \omega^{avg}\right| > x\right] \leq \exp(-I(x)N) \qquad \forall x > 0, n \leq N,$$

$$P(N)\left[\left|\frac{\sum_{n=1}^N \omega_t^n}{N} - \omega^{avg}\right| > x\right] \leq \exp(-I(x)N) \qquad \forall x > 0, t \leq N.$$

Theorem 1. *If $\beta = 1$, then there exists a family, $\{(\Omega^*(N), V^*(N))\}_{N\geq 2}$, of linked VCG mechanisms, one for each N-agent decision problem, such that id is an ε-ex-post equilibrium of each mechanism, and it is common knowledge that*

$$\lim_{N\to\infty} \sup_{\sigma\in\Sigma(N)(id)} P(N)(\sigma(\omega) \neq \omega) = 0,$$

$$\lim_{N\to\infty} \inf_{\sigma\in\Sigma(N)(id)} \frac{\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega)}{\mathbf{E}_{P(N)} S(\boldsymbol{A}(\omega), \omega)} = 1,$$

and

$$\lim_{N\to\infty} \sup_{\sigma\in\Sigma(N)(id)} \frac{\mathbf{E}_{P(N)} C(V^* \circ [\sigma(\omega)])}{\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega)} = 0.$$

Constructing the Linked VCG Mechanisms

Given N and a direct mechanism of the N-model, $\overline{R}(N) := N^2 \overline{v} f'(0)$ is an upper bound on $R_t^n(d|_{t-1}, \omega^n)$ for all $n, t \leq N$.

It is possible to choose a family of positive reals, $\{x(N)\}_{N\geq 2}$ such that

$$\exp(-I(x(N))N) \leq \frac{\frac{\varepsilon}{N}\varepsilon}{N\overline{R}(N)} \quad \forall N,$$
$$\lim_{N\to\infty} x(N) = 0.$$

Given such a family of reals, $\{x(N)\}_{N\geq 2}$, define, for each N,

$$\Omega^{*n}(N) = \left\{\omega^n \in \Omega^n(N) \;\middle|\; \omega_t^n \leq \omega^{max} \ \forall t \leq N, \ \left|\frac{\sum_{t=1}^N \omega_t^n}{N} - \omega^{avg}\right| \leq x(N)\right\}$$

for all $n \leq N$.

This yields a family of stylized type spaces $\{\Omega^*(N)\}_{N\geq 2}$, and, consequently, a family of linked VCG mechanisms $\{(\Omega^*(N), V^*(N))\}_{N\geq 2}$.

Proposition 1 now implies that id is an ε-ex-post equilibrium in each of these linked VCG mechanisms and it is common knowledge that

$$P(N)\left(\sigma(\omega) \neq \omega\right) < \frac{\varepsilon}{N} \quad \forall \sigma \in \Sigma(N)(id).$$

Thus, $\lim_{N\to\infty} \sup_{\sigma\in\Sigma(N)(id)} P(N)(\sigma(\omega) \neq \omega) \leq \lim_{N\to\infty} \frac{\varepsilon}{N} = 0$.

Efficiency

It is possible to choose a family of positive reals, $\{y(N)\}_{N\geq 2}$ such that

$$\lim_{N\to\infty} (N-1)N \exp(-I(y(N))N) = 0,$$
$$\lim_{N\to\infty} y(N) = 0.$$

Define

$$\Omega^{**}(N) = \Bigg\{ \omega \;\; s.t. \;\; \left| \frac{\sum_{n=1}^{N} \omega_t^n}{N} - \omega^{avg} \right| \leq y(N) \; \forall t \leq N,$$
$$\nexists n, t \leq N \;\, P^n(N)(\omega^n \notin \Omega^{*n} \mid \omega^n|_t) > \frac{\varepsilon}{\overline{R}(N)} \Bigg\}$$

By design, $\sigma(\omega) = \omega$ for all $\sigma \in \Sigma(N)(id)$ and $\omega \in \Omega^{**}(N)$, and it is common knowledge that $\lim_{N\to\infty} P(N)(\Omega^{**}(N)) \geq \lim_{N\to\infty} 1 - N\exp(-I(y(N))N) - \frac{\varepsilon}{N} = 1$.

Also,

$$\inf_{\omega\in\Omega^{**}(N)} S(\boldsymbol{A}(\omega), \omega) \geq N^2(\omega^{avg} - y(N)) f\left(\frac{\overline{v}}{\omega^{avg} - y(N)}\right) \geq N^2 \overline{v} f'\left(\frac{\overline{v}}{\omega^{avg}}\right)$$
$$\sup_{\omega\in\Omega(N)} S(\boldsymbol{A}(\omega), \omega) \leq N^2 \overline{v} f'(0)$$

Thus, for $\sigma \in \Sigma(N)(id)$, we have

$$\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega) \geq \mathbf{E}_{P(N)} S(\boldsymbol{A}(\omega), \omega) - (1 - P(\Omega^{**})) N^2 \overline{v} f'(0).$$

Therefore,

$$\lim_{N\to\infty} \sup_{\sigma\in\Sigma(N)(id)} \frac{\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega)}{\mathbf{E}_{P(N)} S(\boldsymbol{A}(\omega), \omega)} \geq \lim_{N\to\infty} 1 - \frac{(1 - P(\Omega^{**})) N^2 \overline{v} f'(0)}{N^2 \overline{v} f'\left(\frac{\overline{v}}{\omega^{avg}}\right)} = 1.$$

Since, by definition of $\boldsymbol{A}$, we have $S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega) \leq S(\boldsymbol{A}(\omega), \omega)$ for all $\omega \in$

$\Omega(N)$, therefore

$$\lim_{N\to\infty} \sup_{\sigma\in\Sigma(N)(id)} \frac{\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega)}{\mathbf{E}_{P(N)} S(\boldsymbol{A}(\omega), \omega)} = 1.$$

Expected Cost

Given $\Omega^*(N)$, $\omega \in \Omega^*(N)$, and $\hat{\omega}^n \in \Omega^{*n}(N)$, we have, by the concavity of f,

$$\begin{aligned}
V^n(\omega) - V^n(\omega^{-n}, \hat{\omega}^n) &= \sum_{t=1}^{N} \Bigg[\sum_{m\neq n} \omega_t^m f\left(\frac{N\overline{v}}{\omega_t^n + \sum_{m\neq n}\omega_t^m}\right) \\
&\qquad - \sum_{m\neq n} \omega_t^m f\left(\frac{N\overline{v}}{\hat{\omega}_t^n + \sum_{m\neq n}\omega_t^m}\right)\Bigg] \\
&\leq \sum_{t=1}^{N} \Bigg[f'\left(\frac{N\overline{v}}{\hat{\omega}_t^n + \sum_{m\neq n}\omega_t^m}\right)\left(\sum_{m\neq n}\omega_t^m\right) \\
&\qquad \cdot \left(\frac{N\overline{v}}{\omega_t^n + \sum_{m\neq n}\omega_t^m} - \frac{N\overline{v}}{\hat{\omega}_t^n + \sum_{m\neq n}\omega_t^m}\right)\Bigg] \\
&= \sum_{t=1}^{N} \Bigg[f'\left(\frac{N\overline{v}}{\hat{\omega}_t^n + \sum_{m\neq n}\omega_t^m}\right)\left(\sum_{m\neq n}\omega_t^m\right) \\
&\qquad \cdot \left(\frac{(\hat{\omega}_t^n - \omega_t^n)N\overline{v}}{\left(\omega_t^n + \sum_{m\neq n}\omega_t^m\right)\left(\hat{\omega}_t^n + \sum_{m\neq n}\omega_t^m\right)}\right)\Bigg].
\end{aligned}$$

This expression then implies

Lemma 3. *Given* $(\Omega^*(N), V^*(N))$ *and* $\omega \in \Omega^{**}(N)$*, we have*

$$V^{*n}(\omega) \leq N2x(N)\overline{v} f'\left(\frac{\overline{v}}{\omega^{avg} + y(N)}\right)\frac{\omega^{avg} + y(N)}{(\omega^{avg} - y(N))^2}$$

Applying the lemma yields

$$\begin{aligned}
\mathbf{E}_{P(N)} C(V^* \circ [\sigma(\omega)]) &\leq N^2 2x(N)\overline{v} f'\left(\frac{\overline{v}}{\omega^{avg} + y(N)}\right)\frac{\omega^{avg} + y(N)}{(\omega^{avg} - y(N))^2} \\
&\quad + \left[N\exp(-I(y(N))N) + \frac{\varepsilon(N)}{N}\right](N-1)N^2\overline{v} f'(0).
\end{aligned}$$

Also,

$$\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega) \geq P(N)(\Omega^{**}(N)) N^2 \overline{v} f' \left(\frac{\overline{v}}{\omega^{avg}} \right)$$

Putting everything together and the theorem is proved:

$$\begin{aligned}
&\lim_{N\to\infty} \sup_{\sigma \in \Sigma(N)(id)} \frac{\mathbf{E}_{P(N)} C(V^* \circ [\sigma(\omega)])}{\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega)} \\
&\leq \lim_{N\to\infty} \frac{2x(N) f'\left(\frac{\overline{v}}{\omega^{avg}+y(N)}\right) \frac{\omega^{avg}+y(N)}{(\omega^{avg}-y(N))^2} + \left[(N-1)N \exp(-I(y(N))N) + \frac{\varepsilon(N)(N-1)}{N}\right] f'(0)}{P(N)(\Omega^{**}(N)) f'\left(\frac{\overline{v}}{\omega^{avg}}\right)} \\
&= 0.
\end{aligned}$$

References

[1] Azevedo, E. M. and E. Budish (2019) "Strategy-proofness in the Large," *Review of Economic Studies* Vol. 86, pp. 81-116

[2] Bergemann, D. and J. Välimäki (2010) "The Dynamic Pivot Mechanism," *Econometrica* Vol. 78, pp. 771-789

[3] Frankel, A. (2014) "Aligned Delegation," *American Economic Review* Vol. 104, pp. 66-83

[4] Green, J. and J-J. Laffont (1977) "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," *Econometrica* Vol. 45, pp. 727-738

[5] Holmström, B. (1979) "Groves' Scheme on Restricted Domains," *Econometrica* Vol. 47, pp. 1137-1144

[6] Jackson, M. O. and A. M. Manelli (1997) "Approximately Competitive Equilibria in Large Finite Economies," *Journal of Economic Theory* Vol. 77, pp. 354-376

[7] Jackson, M. O. and H. F. Sonnenschein (2007) "Overcoming Incentive Constraints by Linking Decisions," *Econometrica* Vol. 75, pp. 241-257

[8] Lee, S. (2017) "Incentive Compatibility of Large Centralized Matching Markets," *Review of Economic Studies* Vol. 84, pp. 444-463

[9] Mailath, G. J., Postlewaite, A., and L. Samuelson (2005) "Contemporaneous Perfect Epsilon-Equilibria," *Games and Economic Behavior* Vol. 53, pp. 126-140

[10] Smet, A., Lund S., and W. Schaninger (2016) "Organizing for the Future," Retrieved from https://www.mckinsey.com/business-functions/organization/our-insights/organizing-for-the-future