Implementation with Statistics

John Y. Zhu[1]

April 4, 2022

Abstract

A method of implementation is introduced for collective decision problems when only some statistics about the type space Ω are known: First, use those statistics to whittle Ω down to a high probability event Ω^*. Then, design a mechanism M^* to ex-post implement the desired outcome treating Ω^* as the type space. Viewed as a mechanism over the true type space Ω, M^* is typically not ex-post. However, under a weaker solution concept I call *ε-ex-post equilibrium*, M^* implements the desired outcome in a high probability subevent of Ω^*. An application to a repeated allocation problem shows how implementation with statistics can yield significantly better results than ex-post implementation.

JEL Codes: D47, D61, D71, D82
Keywords: implementation, mechanism design, statistics, ex-post, resource allocation, social choice, VCG, linking mechanism

[1]johnzhuyiran@ku.edu, Department of Economics, University of Kansas. I thank seminar participants at Texas A&M, National University of Singapore, University of Toronto, and Arizona State University for their advice and comments.

1 Introduction

In a typical collective decision problem involving a mechanism designer and a group 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 statistics about Ω seems quite plausible. Such knowledge could, for example, be extracted from data generated by previous related decision problems or interactions with similar players.

In such a scenario, Bayesian implementation may not be feasible. Designing a mechanism that ex-post implements the desired outcome is an option, but that would involve ignoring the known statistics about Ω, which, intuitively, could be quite sub-optimal: For example, suppose Ω is the set of positive reals, and it is known that the expected value of the type is 1. This single statistic implies that there is a less than 1% chance the realized type exceeds 100. Or, suppose it is known that the type consists of many independent draws from a distribution – a common situation in dynamic decision problems. Then the Law of Large Numbers implies that the realized type is highly likely to lie in a tiny sliver of the type space.

In this paper, I present an approach to implementation, related to the ex-post approach, that can incorporate those potentially valuable statistics: The mechanism designer begins by whittling the true type space Ω down to an event, Ω^*, that the common knowledge statistics imply is of high probability. In the examples considered earlier, Ω^* could be $[0, 100]$ or the tiny sliver implied by the Law of Large Numbers. Next, she designs a direct mechanism M^* to ex-post implement the desired outcome treating Ω^* as the type space. Composing M^* with a retraction $[\cdot] : \Omega \to \Omega^*$ yields a direct mechanism $M^* \circ [\cdot]$ over the true type space Ω that I call a *statistical mechanism.* The mechanism designer then uses this statistical mechanism on the agents.

In what sense does using a statistical mechanism "work?" In the paper, I define what it means for truth-telling to be an *ε-ex-post equilibrium* of a statistical mechanism. ε-ex-post equilibrium is a slight weakening of 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. I then show that, if it is common knowledge that Ω^* is of sufficiently high probability, then truth-telling is an ε-ex-post equilibrium of $M^* \circ [\cdot]$, and, consequently, the mechanism designer can expect with high probability that all agents report the truth at all dates when faced with $M^* \circ [\cdot]$. Moreover, recall, M^* ex-post implements the desired outcome over Ω^*. Putting these two facts together leads us to conclude that, if it is known that Ω^* is of sufficiently high probability, then $M^* \circ [\cdot]$ implements the desired outcome in a high probability subevent of Ω^*.

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

In the second half of the paper, I demonstrate this by comparing the two approaches in the context of a repeated resource allocation problem. The setting is quasilinear, agents have private values and 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 efficient 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 there is common knowledge of some "weak Law of Large Numbers style" statistics, then, in the limit, the mechanism designer can – via a statistical mechanism I call the *linked VCG mechanism* – implement the efficient allocation almost surely at a cost-to-surplus ratio of zero.

The concept of ε-ex-post equilibrium is related to notions of approximate strategy-proofness recently developed by Lee (2017) and Azevedo and Budish (2019). Also related are the contemporaneous perfect ε-equilibrium of Mailath, Postlewaite, and Samuelson (2005) and the dynamic ex-post implementation concept of Bergemann and Välimäki (2002). On the other hand, implementation with statistics is quite distinct from virtual implementation even though both approaches involve the idea of implementation with high probability. See, for example, Abreu and Matsushima (1992). In virtual implementation, the high probability requirement is imposed at the ex-post rather than ex-ante stage – that is, for each type, virtual implementation demands that the mechanism implements the desired outcome with high probability. One can view virtual implementation and implementation with statistics as two orthogonal departures from ex-post implementation. In principle, one could even combine the two methods of implementation (although this is not explored in the current paper): First, whittle an Ω^* as in implementation with statistics, then virtually implement the desired outcome treating Ω^* as the type space.

In the application to repeated resource allocation, my study of linked VCG mechanisms is related to the work of Holmström (1979) on efficient ex-post mechanisms over restricted type spaces. 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 resembles a budget mechanism, revealing a surprising connection between VCG mechanisms and budget mechanisms. A number of papers have shown how budget mechanisms can align incentives across multiple decision problems when transfers are unavailable. The linking mechanism of Jackson and Sonnenschein (2007) is one such budget mechanism, and it is explicitly designed for repeated decision problems like the resource allocation one considered here. See also Frankel (2014). The linking mechanism works in the large numbers limit when the decision problems across dates are iid with known distribution, implementing the efficient allocation almost surely for free. In contrast, the linked VCG mechanism is not free, but, as I will show, it can be designed to work across a wide range of statistical settings, including those that are far from the iid large numbers limit.

2 A Statistical Approach to Implementation

2.1 Decision Problems

Given integers $N \geq 2$ and $T \geq 1$, 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 Ω_t^n is the finite set of date t types for agent n. $D = \Pi_{1 \leq t \leq T} D_t$ is the space of decision sequences, where D_t is the finite set of date t decisions. $U = \Pi_{1 \leq n \leq N} U^n$ is the profile of payoff functions, where $U^n : D \times \Omega^n \to \mathbb{R}$ is agent n's private values payoff function, defined to be

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

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

A *credal set* is a nonempty set of probabilities over Ω. Given a credal set, $\mathcal{P}$, 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 that satisfy some commonly known statistics. The only restriction I impose on $\mathcal{P}$ is that it consists only of *private* probabilities with full support. A probability $\hat{P}$ is private if $\hat{P}(A^n \times \Omega^{-n} |\ \omega|_t) = \hat{P}^n(A^n |\ \omega^n|_t)$ for all n and $A^n \subset \Omega^n$. This restriction implies it is common knowledge that, at each date t, the distribution of agent n's future types is independent of other agents' type histories conditional on agent n's type history.

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

2.2 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 all agents and all dates. If an object $\cdot^n$ is only indexed by agents, then let $\cdot$ denote the profile of $\cdot^n$ across all 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.

2.3 Statistical Mechanisms

Let A and B be two sets of sequences. A map $f : A \to B$ is adapted is $a|_t = a'|_t \Rightarrow f(a)|_t = f(a')|_t$. Given a nonempty $\Omega^{*n} \subset \Omega^n$ for each agent 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}.*

Let $\Omega^* = \Pi_{1\leq n\leq N}\Omega^{*n}$. A direct mechanism over Ω^* is an adapted map $M^* : \Omega^* \to D$. A direct mechanism M^* over Ω^* is ex-post if

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

A *statistical mechanism* is an ex-post direct mechanism M^* over Ω^* composed with a retraction profile $[\cdot] : \Omega \to \Omega^*$. It is a direct mechanism over Ω that is typically not ex-post.

Given a direct mechanism over Ω, an agent n strategy, σ^n, consists of a sequence of maps $\sigma_t^n : D|_{t-1} \times \Omega^n|_t \to \Omega_t^n$. A strategy profile, σ, can be viewed as an adapted map $\sigma : \Omega \to \Omega$. Let id be the truth-telling strategy profile, which induces the identity map on Ω.

2.4 ε-Ex-Post Equilibrium

For the rest of the paper, fix an $\varepsilon > 0$, to be interpreted as "small." In this section, I define what it means for id to be an ε-ex-post equilibrium of a statistical mechanism.

Fix a statistical mechanism, $M^* \circ [\cdot]$, corresponding to some Ω^*. Given $\omega \in \Omega$, agent n's date t regret from continuing to report the truth, assuming he has reported truthfully up through date $t-1$ and assuming all other agents are truth-telling is

$$\max_{\hat{\omega}^n \in \Omega^n \ \ s.t.\ \hat{\omega}^n|_{t-1} = \omega^n|_{t-1}} \sum_{s=t}^{T} \beta^{s-t} \left[u_s^n(M^* \circ [\omega^{-n}, \hat{\omega}^n]|_s, \omega^n|_s) - u_s^n(M^* \circ [\omega]|_s, \omega^n|_s) \right].$$

Fix an upper bound $\overline{R}$ on this regret across all n, t, and ω.

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

$$P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_s) \cdot \overline{R} < \varepsilon \ \ \forall s \leq t \Rightarrow \sigma_t^n(d|_{t-1}, \omega^n|_t) = \omega_t^n. \tag{1}$$

Is this definition of reasonable strategies "reasonable?" Fix a date t and assume the left hand side of (1) is satisfied. By induction, agent n has reported the truth up through date $t-1$. Thus, given his conjecture that other agents are truth-telling, he believes his date t regret from continuing to report the truth is bounded above by $\overline{R}$. Moreover, given that $M^* \circ [\cdot]$ is a statistical mechanism corresponding to Ω^*, he believes he will experience positive regret from continuing to report the truth only if $\omega^n \notin \Omega^{*n}$. Thus, agent n believes that continuing to report the truth is within ε of being ex-post optimal, in which case, it is "reasonable" to assume agent n reports the truth at date t.

One thing worth commenting on about the left hand side of (1): When $t > 1$, agent n observes a nontrivial decision history $d|_{t-1}$, which is informative of other agents' type histories, $\omega^{-n}|_{t-1}$. Since $\omega^{-n}|_{t-1}$ can be informative of ω^n, agent n should, in principle, also condition on $d|_{t-1}$ when formulating his conditional belief about $\omega^n \notin \Omega^{*n}$. This is problematic since, in general, agent n need not know P. Fortunately, the fact that agent n knows P is a private probability and observes his own type history implies that his conditional belief about $\omega^n \notin \Omega^{*n}$ is independent of the other agents' type histories, and, therefore, $d|_{t-1}$.[2]

Let $E^n \subset \Omega$ denote the event $P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_t) \cdot \overline{R} < \varepsilon$ for all $t \leq T$. It is the event in which agent n reports the truth at all dates under any reasonable strategy.

Definition. *id is an ε-ex-post equilibrium if, for each agent n, it is common knowledge $P(\cap_{m \neq n} E^m) \geq 1 - \varepsilon$.*

When *id* is an ε-ex-post equilibrium, I will interpret it to mean that it is common knowledge each agent n will play a reasonable strategy. In this case, the definition of ε-ex-post equilibrium implies that it is common knowledge each agent's conjecture that all other agents are truth-telling is, at worst, almost correct.

In a typical application, the set of reasonable agent n strategies will not be known to anyone except agent n. Nevertheless, one will still be able to deduce things just from the knowledge that each agent n plays a reasonable strategy. In particular, it is common knowledge that the probability all agents report the truth at all dates is at least $1 - 2\varepsilon$.

ε-ex-post equilibrium generalizes ex-post equilibrium in the following sense: If *id* is an ex-post equilibrium of $M^* \circ [\cdot]$ then $\overline{R}$ can be set to 0. If $\overline{R} = 0$ then $E^n = \Omega$ for all n. In this case, *id* is obviously an ε-ex-post equilibrium of $M^* \circ [\cdot]$.

It is also worth comparing the definition of ε-ex-post equilibrium with that of perfect Bayesian equilibrium. In a perfect Bayesian equilibrium,

I. Each agent's beliefs are "reasonable" given his strategy and his conjecture about the other agents' strategies (Bayes' Rule).

II. Each agent's strategy is "reasonable" given his beliefs and his conjecture about the other agents' strategies (sequential rationality).

III. Each agent correctly conjectures the other agents' strategies.

In terms of ε-ex-post equilibrium, the relevant beliefs agent n must form are those about the probability his type will land outside Ω^{*n}. Since agent n knows his own marginal, P^n, these beliefs are $\{P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_t)\}_{t=1}^T$ which, by definition, satisfy Bayes' Rule. The definition of what it means for σ^n to be reasonable corresponds to the sequential rationality condition.

Finally, the condition that it is common knowledge each agent's conjecture that all other agents are truth-telling is, at worst, almost correct corresponds to (a slight

[2]Alternatively, we could just assume that agent n does not observe $d|_{t-1}$ or is unable to infer anything from observing $d|_{t-1}$, perhaps because it is too mentally taxing to make such inferences. In either case, we can drop the restriction that credal sets only contain private probabilities.

weakening of) the correct conjectures condition. This almost correct condition is an ex-ante one. It is possible that, late in a reasonable strategy profile, the continuation strategies of all other agents are far from truth-telling. Our theory of reasonable strategies can accommodate agents possibly "updating" their conjectures.

"Updating" a conjecture is not a well-defined, simple, mechanical procedure like Bayesian updating. Whatever one imagines "updating" to be, it will likely involve some costly information gathering and costly contemplation on the part of the agent. Thus, it stands to reason that an agent will attempt to update his conjecture only if he believes the benefit of an updated conjecture is not too small. But his belief about the benefit of an updated conjecture must be computed under his current conjecture. Now, let us revisit (1). Recall, if the left hand side is satisfied at some date t, then, assuming he has not updated his conjecture up through date $t-1$, agent n believes at date t that continuing with truth-telling is almost ex-post optimal. This implies he believes there is little benefit to updating his conjecture at date t. In this case, we can modify the right hand side of (1) so that not only does he report the truth at date t, but he also does not update his conjecture at date t. As a result, the set of reasonable strategy profiles is unchanged. The event in which an agent n reports the truth at all dates under any reasonable strategy remains E^n, and the definition of what it means for id to be an ε-ex-post equilibrium can be left unchanged.

2.5 Implementation with Statistics

Fix a statistical mechanism $M^* \circ [\cdot]$ corresponding to some Ω^* and an upper bound $\overline{R}$ on regret across all n, t, and ω.

A *desired outcome* is a nonempty-valued correspondence $DO : \Omega \to D$. For example, in an auction decision problem, a decision would be an allocation of the object along with payments from the bidders, and, for a mechanism designer who desires efficiency, $DO(\omega)$ could be the set of decisions that involve allocating the object to the bidder with the highest valuation.

We say M^* ex-post implements the desired outcome over Ω^* if $M^*(\omega) \in DO(\omega)$ for all $\omega \in \Omega^*$. We say $M^* \circ [\cdot]$ implements the desired outcome with probability at least p if id is an ε-ex-post equilibrium and it is common knowledge that $P(M^* \circ [\sigma(\omega)] \in DO(\omega)) \geq p$ for all reasonable σ.

Theorem 1. *Suppose there exists a $c \in [0, \varepsilon]$ such that it is common knowledge $P^n(\Omega^{*n}) \geq 1 - \frac{c}{N-1} \cdot (\frac{\varepsilon}{\overline{R}} \wedge 1)$ for all agents n, and suppose M^* ex-post implements the desired outcome over Ω^*. Then $M^* \circ [\cdot]$ implements the desired outcome with probability at least $1 - \frac{Nc}{N-1}$.*

Theorem 1 formalizes the idea that, if it is common knowledge that Ω^* is of sufficiently high probability and M^* ex-post implements the desired outcome over Ω^*, then, over the true type space Ω, $M^* \circ [\cdot]$ implements the desired outcome with high probability.

3 Application: Repeated Resource Allocation

3.1 Model

A principal 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, $f : [0, \infty) \to [0, \infty)$. f is a strictly concave, C^1 function 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 efficient allocation of her resource each date.

This model can be generalized to allow for many different kinds of resources and project types. In the generalization, the principal possesses a vector quantity of resources, with each entry corresponding to a different kind of resource. A project type is a strictly concave, C^1 function over some subset of resources. Everything I am about to prove extends to this more general setting.

One application of the 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. 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. If labor is specialized, then each specialization corresponds to a different kind of resource. Projects are departmental tasks that map labor (or vectors of specialized labor) to output. Transfers from the principal to agents correspond to incentive pay for department managers.

3.2 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
- $U^n((a, w), \omega^n) = \sum_{t=1}^T \beta^{t-1} \omega_t^n f\left(\frac{a_t^n}{\omega_t^n}\right) + \beta^{T-1} w^n$ for all n.

Technically speaking, Ω and D should be finite sets. However, the theory developed earlier naturally extends to handle such infinite Ω and D.

In addition, define the following auxiliary quantities,

- agent n's contribution to surplus: $S^n(a^n, \omega^n) = \sum_{t=1}^T \beta^{t-1} \omega_t^n f\left(\frac{a_t^n}{\omega_t^n}\right)$,
- surplus: $S(a, \omega) = \sum_{n=1}^N S^n(a^n, \omega^n)$, and
- cost: $C(w) = \sum w^n$.

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

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

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 desired outcome is $DO(\omega) = \{(\boldsymbol{A}(\omega), w) \mid w^n \geq 0 \ \forall n\}$ for all $\omega \in \Omega$.

I now compare the ex-post and statistical approaches to implementation, with an emphasis on the cost of implementation.

3.3 The Unlinked VCG Mechanism

Suppose the principal wants to ex-post implement the efficient allocation. One option is to run a separate Vickrey-Clarke-Groves (VCG) mechanism each date, paying each agent the sum of all other agents' contributions to surplus, quoted in date T terms:

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

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

Proposition 1. *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 1 is a corollary of Proposition 2 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

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

As the number of agents tends to infinity, so does the cost-to-surplus ratio.

3.4 The Linked VCG Mechanism

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

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

Since the VCG mechanism is ex-post, so is the linked VCG mechanism. In fact,

Proposition 2. *If Ω^* is smoothly path-connected, then the linked VCG mechanism $(\boldsymbol{A}|_{\Omega^*}, V^*)$ is the cheapest efficient ex-post direct mechanism over Ω^*.*

Proposition 2 is a novel application of Theorem 1 of Holmström (1979) about the necessity of VCG mechanisms over restricted type spaces. The novelty derives from the kind of type space restriction being considered. Holmström (1979) focuses on preference driven restrictions of the type space. In a canonical example, the type space is the space of all continuous utility functions, and the restriction is to the subspace of all concave utility functions. In contrast, I focus on statistically driven restrictions of the type space, where the type space is whittled to an event commonly known to be of high probability. Either way, the underlying mathematics – essentially the envelope theorem – is the same.

Theorem 1 implies that any statistical mechanism corresponding to $(\boldsymbol{A}|_{\Omega^*}, V^*)$ implements the efficient allocation with high probability if it is common knowledge Ω^{*n} is of sufficiently high probability for each agent n. As an abuse of nomenclature, call any statistical mechanism corresponding to $(\boldsymbol{A}|_{\Omega^*}, V^*)$ 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 "weak Law of Large Numbers style" statistics, then, by taking the statistical 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.

In order to prove such a result, I now introduce a family of decision problems parameterized by the number of agents/dates. This will allow me to take the limit as the number of agents/dates goes to infinity.

3.5 A Family of Decision Problems

Fix a quantity $q > 0$ of the resource and a project type f. Introduce a family of decision problems parameterized by the number of agents N, appending (N) to parameters of the N-agent decision problem. Since we are assuming agents are patient, set $\beta(N) = 1$ for all N. In addition, let us assume the project type stays the same across decision problems and the amount of available resources scales proportionally with the number of agents: $(f(N), \overline{q}(N)) = (f, Nq)$. Also, since we want the number of agents and dates to go to infinity together, let us assume $T(N) = N$.

Lastly, the following assumption about the family of credal sets captures the idea that there is common knowledge of some "weak Law of Large Numbers style" statistics:

Assumption 1. *There exist*

- *a bounded set of positive reals* $\{\omega^{avg}\} \cup \{\omega^{avg,n}\}_{n=1}^{\infty}$,
- *weakly decreasing functions* $g, G : [0,\infty) \to [0,\infty)$ *satisfying* $\lim_{z\to\infty} z^5 g(z) = \lim_{z\to\infty} z^3 G(z) = 0$,
- *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

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

If it is common knowledge that the array of endowments $\{\omega_t^n\}_{1\leq n,t<\infty}$ is iid with mean μ and variance σ^2, then the Central Limit Theorem implies Assumption 1 is satisfied. However, Assumption 1 can accommodate significant departures from the baseline iid case, along both the identically distributed and the independent dimensions. For example, suppose it is common knowledge that the array of endowments is independent but the only distributional assumption is that it is common knowledge they share a mean μ and upper bound $\overline{\omega}$. Then Hoeffding's Inequality says that, for each N,

$$P(N)\left[\left|\frac{\sum_{s\neq t}\omega_s^n}{N-1} - \mu\right| > x\right] \leq 2\exp\left(-\frac{2(N-1)x^2}{\overline{w}^2}\right),$$
$$P(N)\left[\left|\frac{\sum_{m\neq n}\omega_t^m}{N-1} - \mu\right| > x\right] \leq 2\exp\left(-\frac{2(N-1)x^2}{\overline{w}^2}\right),$$

for all $x > 0$ and $t, n \leq N$. By defining $g(z) = G(z) = 2\exp(-z)$ and $I(x) = \frac{2x^2}{\overline{\omega}^2}$, we see that Hoeffding's Inequality implies Assumption 1 is satisfied. It is also clear that Assumption 1 allows for the possibility of significant correlation between endowments that differ in both the agent and time dimensions. This means many pairs of endowments can be highly correlated while still satisfying Assumption 1.

Theorem 2. *There exists a family, $\{(\Omega^*(N), V^*(N))\}_{N\geq 2}$, of linked VCG mechanisms, along with a family, $\{\overline{R}(N)\}_{N\geq 2}$, of upper bounds on regret, such that id is an*

ε-ex-post equilibrium of each statistical mechanism, and it is common knowledge that

$$\lim_{N\to\infty} \inf_{\text{reasonable } \sigma} P(N)\left(\omega \in \Omega^*(N),\ \sigma(\omega) = \omega\right) = 1,$$

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

and

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

Theorem 2 says that, in the limit, the linked VCG mechanism can implement the efficient allocation almost surely, generating an expected surplus to expected efficient surplus ratio of 1, at an expected cost to expected surplus ratio of 0.

Let us gain some intuition for this result. When N is large, the efficient allocation gives each date t project an amount of resources equal to

$$\frac{Nq}{\sum_{m=1}^{N} \omega_t^m} \approx \frac{q}{\omega^{avg}}.$$

Thus, the efficient marginal productivity of resource is approximately always $f' := f'(q/\omega^{avg})$. When N is large, we can approximately whittle each agent n's type space to

$$\Omega^{*n}(N) \approx \left\{ \omega^n \;\middle|\; \frac{\sum_{t=1}^{N} \omega_t^n}{N} = \omega^{avg,n} \right\},$$

while still ensuring that the probability $\omega \in \Omega^*(N)$ and $\sigma(\omega) = \omega$ under any reasonable σ is high. Fix such an ω. Let us now approximate agent n's transfer under the linked VCG mechanism $(\Omega^*(N), V^*(N))$.

$$\begin{aligned}
V^{*n}(N)(\omega) &= V^n(N)(\omega) - \inf_{\hat{\omega}^n \in \Omega^{*n}(N)} V^n(N)(\omega^{-n}, \hat{\omega}^n) \\
&= \sup_{\hat{\omega}^n \in \Omega^{*n}(N)} \left(V^n(N)(\omega^{-n}, 0) - V^n(N)(\omega^{-n}, \hat{\omega}^n)\right) \\
&\qquad - \left(V^n(N)(\omega^{-n}, 0) - V^n(N)(\omega)\right) \\
&\approx \sup_{\hat{\omega}^n \in \Omega^{*n}(N)} \left(\sum_{t=1}^{N} \hat{\omega}_t^n\right) \cdot \frac{q}{\omega^{avg}} \cdot f' - \left(\sum_{t=1}^{N} \omega_t^n\right) \cdot \frac{q}{\omega^{avg}} \cdot f' \\
&\approx N\omega^{avg,n} \cdot \frac{q}{\omega^{avg}} \cdot f' - \left(\sum_{t=1}^{N} \omega_t^n\right) \cdot \frac{q}{\omega^{avg}} \cdot f' \qquad (2)
\end{aligned}$$

$$\approx N\omega^{avg,n} \cdot \frac{q}{\omega^{avg}} \cdot f' - N\omega^{avg,n} \cdot \frac{q}{\omega^{avg}} \cdot f' \approx 0 \text{ fraction of } N.$$

Summing over all agents yields a cost that is approximately a zero fraction of N^2. Since expected surplus is on the order of N^2, the expected cost to expected surplus ratio is approximately zero.

As equation (2) in the derivation of agent n's approximate transfer makes clear, in the limit, the linked VCG mechanism resembles a budget mechanism where each agent n is given a budget of $N\omega^{avg,n} \cdot \frac{q}{\omega^{avg}} \cdot f'$ and the price of the resource is set to f' each date. The agent is then free to choose how much resources to buy each date subject to his budget constraint, keeping any funds he does not use by the end of the model.

Given the connection to budget mechanisms, it is natural to compare the linked VCG mechanism to the linking mechanism of Jackson and Sonnenschein (2007) (hereafter J-S), which is an explicit budget mechanism specifically designed for repeated decision problems like the one considered here. In the context of the repeated resource allocation problem, the linking mechanism of J-S allows each agent n to report whatever endowment process he wants subject to the budget constraint that the empirical distribution of endowments matches the probability distribution of endowments.

When there are many dates, and $\{\omega_t^n\}$ is independent across agents and dates, and identically distributed with known distribution across dates holding the agent fixed, the linking mechanism implements the efficient allocation almost surely for free. In contrast, the best the linked VCG mechanism can do under these probabilistic assumptions is to implement the efficient allocation almost surely at an expected cost that is a vanishingly small fraction of expected surplus – and even this result requires there to be many agents in addition to many dates.

However, the linked VCG mechanism does have some strengths. Our discussion following Assumption 1 implies that Theorem 2 remains valid in settings that significantly relax the probabilistic assumptions of J-S: Many pairs of endowments can be highly correlated; holding the agent fixed, endowments across dates can be far from identically distributed; and the mechanism designer need not know any agent's endowment distribution.

Perhaps the biggest strength of the linked VCG mechanism is that it is not just a single mechanism, perfectly adapted to the large numbers limit. Given any credal set, the mechanism designer can whittle Ω down to a sufficiently high probability Ω^* and create the corresponding linked VCG mechanism. This linked VCG mechanism will implement the efficient allocation with high probability. Moreover, if Ω^* is smoothly path-connected, then it would be the optimal efficient ex-post mechanism if we were to approximate the true type space with Ω^*. A reasonable interpretation of these facts is that the linked VCG mechanism is a family of mechanisms that is adaptable – if not perfectly – to a wide variety of statistical settings. That adaptability ultimately derives from the fact that the linked VCG mechanism is a product of the statistical approach to implementation introduced in this paper.

4 Appendix

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, $[\omega'^n]|_t = \omega'^n|_t = \omega''|_t = [\omega''^n]|_t$. $\square$

Proof of Theorem 1. First, suppose $\overline{R} \leq \varepsilon$. Then the only reasonable strategy profile is id. Thus, for any $\omega \in \Omega^*$, $M^* \circ [\sigma(\omega)] \in DO(\omega)$ for all reasonable σ. Since it is common knowledge $P^n(\Omega^{*n}) \geq 1 - \frac{c}{N-1}$, therefore, by de Morgan's Law, it is common knowledge $P(\Omega^*) \geq 1 - \frac{Nc}{N-1}$ and the result is proved.

Next, suppose $\overline{R} > \varepsilon$. Define $X^n_t(\omega) := P^n(\omega^n \notin \Omega^{*n} \mid \omega^n|_t)$. Extend the sequence by one date by defining $X^n_{T+1} = X^n_T$. It is common knowledge X^n is a nonnegative martingale with expected value $X^n_0 = P^n(\omega^n \notin \Omega^{*n}) \leq \frac{c}{N-1} \cdot \frac{\varepsilon}{\overline{R}}$.

Let τ^n denote the stopping time when X^n_t first weakly exceeds $\frac{\varepsilon}{\overline{R}}$. If X^n_t never weakly exceeds $\frac{\varepsilon}{\overline{R}}$, then set $\tau^n = T+1$. Then $E^n \subset \Omega$ is the event $\tau^n = T+1$. By Doob's optional stopping theorem, we have

$$\frac{c}{N-1} \cdot \frac{\varepsilon}{\overline{R}} \geq X^n_0 = \mathbf{E}X^n_{\tau^n} = \mathbf{E}X^n_{\tau^n} 1_{\tau^n \leq T} + \mathbf{E}X^n_{\tau^n} 1_{\tau^n = T+1} \geq \mathbf{E}X^n_{\tau^n} 1_{\tau^n \leq T}.$$

Since $\mathbf{E}X^n_{\tau^n} 1_{\tau^n \leq T} \geq \frac{\varepsilon}{\overline{R}} \cdot (1 - P(E^n))$, it is common knowledge $P(E^n) \geq 1 - \frac{c}{N-1}$.

By de Morgan's Law, for each agent n, $P(\cap_{m \neq n} E^m) \geq 1 - \sum_{m \neq n}(1 - P(E^m)) \geq 1 - c \geq 1 - \varepsilon$. So id is ε-ex-post equilibrium.

Similarly, $P(\cap_n E^n) \geq 1 - \frac{Nc}{N-1}$. By definition, if $\omega \in \cap_n E^n$ then $\sigma(\omega) = \omega$ for any reasonable σ. Moreover, $\cap_n E^n \subset \Omega^*$. Thus, if $\omega \in \cap_n E^n$ then $M^* \circ [\sigma(\omega)] = M^*(\omega) \in DO(\omega)$. This proves $M^* \circ [\cdot]$ implements the desired outcome with probability at least $1 - \frac{Nc}{N-1}$. $\square$

Proof of Proposition 2. Let Ω^* be smoothly path-connected and let $(\boldsymbol{A}|_{\Omega^*}, W)$ be an efficient ex-post direct mechanism over Ω^*. It suffices to show that $V^{*n}(\omega) \leq W^n(\omega)$ for all n and $\omega \in \Omega^*$.

Suppose not. Then there exists an n and an $\omega \in \Omega^*$ such that $W^n(\omega) = V^{*n}(\omega) - \delta$ for some $\delta > 0$. By definition of V^*, there exists a $\hat{\omega} \in \Omega^*$ satisfying $\hat{\omega}^{-n} = \omega^{-n}$ and $V^{*n}(\hat{\omega}) < \delta$. Theorem 1 of Holmström (1979) implies

$$W^n(\hat{\omega}) - V^{*n}(\hat{\omega}) = W^n(\omega) - V^{*n}(\omega) = -\delta.$$

Thus, $W^n(\hat{\omega}) = W^n(\hat{\omega}) - V^{*n}(\hat{\omega}) + V^{*n}(\hat{\omega}) < -\delta + \delta = 0$. Contradiction. $\square$

4.1 Proof of Theorem 2

I specialize to the case where $\omega^{avg} = \omega^{avg,n}$ for all n. The proof is easily adapted to the general case.

4.1.1 Constructing the Family of Linked VCG Mechanisms

Given N and any direct mechanism of the N-model, we can set $\overline{R}(N) := N^2 q f'(0) \vee \varepsilon$.

Lemma 2. *There exists a family of positive reals, $\{\varepsilon(N), x(N)\}_{N\geq 2}$ such that*

$$\begin{aligned} \varepsilon(N) &\leq \varepsilon && \forall N, \\ Ng(I(x(N))(N-1)) &\leq \frac{\frac{\varepsilon(N)}{N}\varepsilon}{(N-1)\overline{R}(N)} && \forall N, \\ \lim_{N\to\infty} \varepsilon(N) &= \lim_{N\to\infty} x(N) = 0. \end{aligned}$$

Proof. Given integer $k > 0$, since $\lim_{z\to\infty} z^5 g(z) = 0$, we have

$$\lim_{N\to\infty} (N-1)N^2\overline{R}(N) g\left(I\left(\frac{1}{k}\right)(N-1)\right) = 0.$$

Thus, there exists an integer N_k such that for all $N \geq N_k$,

$$(N-1)N^2\overline{R}(N) g\left(I\left(\frac{1}{k}\right)(N-1)\right) \leq \frac{\varepsilon}{k} \cdot \varepsilon.$$

Obviously, the sequence of N_k can be chosen to be strictly increasing.

Since $\lim_{x\to\infty} I(x) = \infty$, there exists an x_0 such that

$$N_1 g(I(x_0)) \leq \frac{\frac{\varepsilon}{N_1}\varepsilon}{(N_1-1)\overline{R}(N_1)}.$$

For $N < N_1$, define $\varepsilon(N) = \varepsilon$ and $x(N) = x_0$. For all integers $k > 0$, and $N \in \{N_k, N_k+1, \ldots N_{k+1}-1\}$, define $\varepsilon(N) = \frac{\varepsilon}{k}$ and $x(N) = \frac{1}{k}$. The lemma is proved. □

Fix a family, $\{\varepsilon(N), x(N)\}_{N\geq 2}$, as in Lemma 2, and define, for each N,

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

for all $n \leq N$.

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

Assumption 1 and de Morgan's Law imply it is common knowledge that

$$P^n(N)(\Omega^{*n}(N)) \geq 1 - \frac{\frac{\varepsilon(N)}{N}\varepsilon}{(N-1)\overline{R}(N)}.$$

Theorem 1 then implies that *id* is an ε-ex-post equilibrium of each statistical mechanism and it is common knowledge that

$$P(N)\left(\omega \in \Omega^*(N),\ \sigma(\omega) = \omega \quad \forall \text{ reasonable } \sigma\right) > 1 - \frac{\varepsilon(N)}{N-1}.$$

This implies the first part of Theorem 2.

4.1.2 Expected Surplus

Lemma 3. *There exists a family of positive reals, $\{y(N)\}_{N\geq 2}$, such that $y(N) < \omega^{avg}$ for all N, and*

$$\lim_{N\to\infty} (N-1)N^2 G(I(y(N))(N-1)) = 0,$$
$$\lim_{N\to\infty} y(N) = 0.$$

Proof. Fix an integer k_0 such that $\frac{1}{k_0} < \omega^{avg}$. Since $\lim_{z\to\infty} z^3 G(z) = 0$, for each integer $k \geq k_0$, there exists an integer N_k such that for all $N \geq N_k$,

$$(N-1)N^2 G\left(I\left(\frac{1}{k}\right)(N-1)\right) \leq \frac{1}{k}.$$

Choose N_k strictly increasing in k. Fix a $y_0 \in (0, \omega^{avg})$. Then define $y(N) = y_0$ for all $N < N_{k_0}$, and $y(N) = \frac{1}{k}$ for all $N \in \{N_k, N_k+1, \dots N_{k+1}-1\}$, $k \geq k_0$. $\square$

Fix a family, $\{y(N)\}_{N\geq 2}$ as in Lemma 3, and define $\Omega^{**}(N) =$

$$\left\{\omega \;\middle|\; \left|\frac{\sum_{m\neq n}\omega_t^m}{N-1} - \omega^{avg}\right| \leq y(N)\ \forall t, n \leq N,\ \omega \in \Omega^*(N),\ \sigma(\omega) = \omega \quad \forall \text{ reasonable } \sigma\right\}$$

Assumption 1 and de Morgan's Law imply it is common knowledge $P(N)(\Omega^{**}(N)) \geq 1 - N^2 G(I(y(N))(N-1)) - \frac{\varepsilon(N)}{N-1}$. So $\lim_{N\to\infty} P(N)(\Omega^{**}(N)) = 1$.

Lemma 4. *The efficient surpluses, $S(\boldsymbol{A}(\omega), \omega)$, satisfies the following bounds:*

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

Proof. $\omega \in \Omega^{**}(N)$ implies $\left|\frac{\sum_{m\neq n}\omega_t^m}{N-1} - \omega^{avg}\right| \leq y(N)$ for all $n, t \leq N$, which then implies $\left|\frac{\sum_n \omega_t^n}{N} - \omega^{avg}\right| \leq y(N)$ for all $t \leq N$. Thus, the lowest possible surplus is generated when $\sum_n \omega_t^n = N(\omega^{avg} - y(N))$ for all $t \leq N$. In this case, the surplus generated is $N^2(\omega^{avg} - y(N)) f\left(\frac{q}{\omega^{avg}-y(N)}\right)$. On the other hand, the surplus generated is always weakly less than what is generated if $\omega_t^n = \infty$ for all $n, t \leq N$. In this case, the surplus generated is $N^2 q f'(0)$. $\square$

Whenever $\omega \in \Omega^{**}$, all agents report the truth under any reasonable σ. Thus, the second part of Lemma 4 implies, under any reasonable σ,

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

And now, the first part of Lemma 4 implies

$$\begin{aligned}\lim_{N\to\infty} \inf_{\text{reasonable }\sigma} \frac{\mathbf{E}_{P(N)} S(\boldsymbol{A}|_{\Omega^*(N)} \circ [\sigma(\omega)], \omega)}{\mathbf{E}_{P(N)} S(\boldsymbol{A}(\omega), \omega)} &\geq \\ \lim_{N\to\infty} 1 - &\frac{(1 - P(N)(\Omega^{**}(N)))q f'(0)}{P(N)(\Omega^{**}(N))(\omega^{avg} - y(N)) f\left(\frac{\overline{q}}{\omega^{avg}-y(N)}\right)} = 1.\end{aligned}$$

This proves the second part of Theorem 2.

4.1.3 Expected Cost

Lemma 5. *Let $\omega^n \in \Omega^{*n}(N)$. Then $\omega_t^n \leq \omega^{avg} + (2N-1)x(N)$ for all $t \leq N$.*

Proof. $\omega^n \in \Omega^{*n}(N)$ implies $\left|\frac{\sum_{s\neq t}\omega_s^n}{N-1} - \omega^{avg}\right| \leq x(N)$ for all $t \leq N$, which then implies $\left|\frac{\sum_s \omega_s^n}{N} - \omega^{avg}\right| \leq x(N)$ or, equivalently, $\left|\sum_{s\neq t}\omega_s^n + \omega_t^n - N\omega^{avg}\right| \leq Nx(N)$ for all $t \leq N$. Thus,

$$\begin{aligned}\omega_t^n \leq N\omega^{avg} - \sum_{s\neq t}\omega_s^n + Nx(N) &\leq N\omega^{avg} - (N-1)(\omega^{avg} - x(N)) + Nx(N) \\ &= \omega^{avg} + (2N-1)x(N).\end{aligned}$$

$\square$

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

$$\begin{aligned}V^n(N)(\omega) - V^n(N)(\omega^{-n}, \hat{\omega}^n) = \sum_{t=1}^{N}\Bigg[&\sum_{m\neq n}\omega_t^m f\left(\frac{Nq}{\omega_t^n + \sum_{m\neq n}\omega_t^m}\right) \\ &- \sum_{m\neq n}\omega_t^m f\left(\frac{Nq}{\hat{\omega}_t^n + \sum_{m\neq n}\omega_t^m}\right)\Bigg]\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{t=1}^{N}\Bigg[f'\left(\frac{Nq}{\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m}\right)\left(\sum_{m\neq n}\omega_t^m\right)\\
&\qquad\cdot\left(\frac{Nq}{\omega_t^n+\sum_{m\neq n}\omega_t^m}-\frac{Nq}{\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m}\right)\Bigg]\\
&= \sum_{t=1}^{N}\Bigg[f'\left(\frac{Nq}{\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m}\right)\left(\sum_{m\neq n}\omega_t^m\right)\\
&\qquad\cdot\left(\frac{(\hat{\omega}_t^n-\omega_t^n)Nq}{\left(\omega_t^n+\sum_{m\neq n}\omega_t^m\right)\left(\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m\right)}\right)\Bigg].
\end{aligned}\tag{3}$$

This expression then implies

Lemma 6. *Linked VCG transfers are asymptotically dominated by N over $\Omega^{**}(N)$.*

Formally, there exists a family of positive reals $\{a(N), b(N)\}_{N\geq 2}$ satisfying $a(N) < b(N)$ for all N, and $\lim_{N\to\infty} a(N) = \lim_{N\to\infty} b(N) < \infty$, such that for all agents $n \leq N$ and $\omega \in \Omega^{**}(N)$, we have

$$V^{*n}(N)(\omega) \leq N2x(N)a(N) + N(\omega^{avg} + x(N))(b(N) - a(N)).$$

Proof. Consider the quantity

$$f'\left(\frac{Nq}{\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m}\right)\left(\sum_{m\neq n}\omega_t^m\right)\cdot\left(\frac{Nq}{\left(\omega_t^n+\sum_{m\neq n}\omega_t^m\right)\left(\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m\right)}\right).\tag{4}$$

Given $\omega \in \Omega^{**}(N)$, and $\hat{\omega}^n \in \Omega^{*n}(N)$, let us bound from above and below this quantity.

Lemma 5 and the fact that $\omega \in \Omega^{**}(N)$ imply

$$f'\left(\frac{Nq}{\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m}\right)\in\left[f'\left(\frac{\frac{N}{N-1}q}{\omega^{avg}-y(N)}\right), f'\left(\frac{\frac{N}{N-1}q}{\frac{\omega^{avg}}{N-1}+\frac{2N-1}{N-1}x(N)+\omega^{avg}+y(N)}\right)\right].$$

Similarly,

$$\left(\sum_{m\neq n}\omega_t^m\right)\cdot\left(\frac{Nq}{\left(\omega_t^n+\sum_{m\neq n}\omega_t^m\right)\left(\hat{\omega}_t^n+\sum_{m\neq n}\omega_t^m\right)}\right)\in$$
$$\left[\frac{(\omega^{avg}-y(N))q}{(\omega^{avg}+y(N))\left(\frac{\omega^{avg}}{N-1}+\frac{2N-1}{N-1}x(N)+\omega^{avg}+y(N)\right)}, \frac{(\omega^{avg}+y(N))\frac{N}{N-1}q}{(\omega^{avg}-y(N))(\omega^{avg}-y(N))}\right]$$

So, define

$$a(N) := f'\left(\frac{\frac{N}{N-1}q}{\omega^{avg} - y(N)}\right) \cdot \frac{(\omega^{avg} - y(N))q}{(\omega^{avg} + y(N))\left(\frac{\omega^{avg}}{N-1} + \frac{2N-1}{N-1}x(N) + \omega^{avg} + y(N)\right)},$$

$$b(N) := f'\left(\frac{\frac{N}{N-1}q}{\frac{\omega^{avg}}{N-1} + \frac{2N-1}{N-1}x(N) + \omega^{avg} + y(N)}\right) \cdot \frac{(\omega^{avg} + y(N))\frac{N}{N-1}q}{(\omega^{avg} - y(N))(\omega^{avg} - y(N))}.$$

Clearly, $a(N) < b(N)$ for all N. Moreover

$$\lim_{N\to\infty} a(N) = \lim_{N\to\infty} b(N) = \frac{q}{\omega^{avg}} f'\left(\frac{q}{\omega^{avg}}\right).$$

And now, we can bound from above the right hand side of (3) by

$$\begin{aligned}\sum_{t=1}^{N}\left(\hat{\omega}_t^n b(N) - \omega_t^n a(N)\right) &= \left(\sum_{t=1}^{N}\hat{\omega}_t^n - \sum_{t=1}^{N}\omega_t^n\right) a(N) + \sum_{t=1}^{N}\hat{\omega}_t^n\left(b(N) - a(N)\right)\\ &\leq N2x(N)a(N) + N(\omega^{avg} + x(N))\left(b(N) - a(N)\right).\end{aligned}$$

The result now follows from the observation that

$$V^{*n}(N)(\omega) = \sup_{\hat{\omega}^n \in \Omega^{*n}}\left(V^n(N)(\omega) - V^n(N)(\omega^{-n}, \hat{\omega}^n)\right).$$

□

Applying Lemma 6, we have

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

Also,

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

Putting everything together and the third and final part of Theorem 2 is proved:

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

References

[1] Abreu, D. and H. Matsushima (1992) "Virtual Implementation in Iteratively Undominated Strategies: Complete Information," *Econometrica* Vol. 60, pp. 993-1008

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

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

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

[5] 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

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

[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`