Repeated Resource Allocation

John Y. Zhu[1]

August 30, 2021

Abstract

A principal repeatedly allocates resources to a set of agents. Each agent is privately endowed with a sequence of projects that use resources to generate payoffs. Ergodicity implies each agent's average endowment likely lies in some restricted confidence interval. I introduce the *linked VCG mechanism* and show it is the cheapest efficient ex-post mechanism if agents' endowment processes are stylized so that average endowments *must* lie in those confidence intervals.

Linking VCG mechanisms can yield significant cost savings for the principal, demonstrating how trimming the type space of a mechanism design problem by removing elements known to be unlikely can greatly improve the solution. The attractive properties of the linked VCG mechanism are approximately preserved when the stylized restriction on average endowments is relaxed.

JEL Codes:
Keywords:

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

1 Introduction.

A principal has stocks of R kinds of durable, divisible resources. Each date, each agent is privately endowed with some divisible projects. A project lasts one date, using some nonnegative R-vector of resources to generate a unit payoff. There are I different project types, parameterized by the resource input, and an endowment of projects is a nonnegative I-vector. The principal repeatedly allocates her durable resources and makes monetary transfers to the agents depending on the history of reported endowments. Each agent's quasilinear utility is the discounted sum of his project payoffs and transfers from the principal. Agents are protected by limited liability, so transfers must be nonnegative. The principal seeks a mechanism that achieves allocative efficiency while minimizing transfers.

One way to achieve allocative efficiency is to run a Vickrey-Clarke-Groves (VCG) mechanism each date, paying each agent the sum of everyone else's contributions to surplus. But this is expensive for the principal, and does not take advantage of the repeated structure of the problem. Call this the unlinked VCG mechanism.

To explore the possibility of better mechanisms, suppose the time horizon is T dates and it is common knowledge that the T-date average of each agent n's endowment process must lie in some confidence interval, $[\underline{\mu}^n, \overline{\mu}^n]$, where $\underline{\mu}^n \leq \overline{\mu}^n$ are nonnegative I-vectors. Consider the following *linked VCG mechanism*: The principal modifies the unlinked VCG mechanism by requiring the T-date average of each agent n's reported endowments to be in $[\underline{\mu}^n, \overline{\mu}^n]$, and by postponing transfers until the end. Assume common discounting so that postponing transfers is costless. At the end of date T, let y denote the transfer to agent n under the unlinked VCG mechanism. Let x denote the minimum transfer to agent n under the unlinked VCG mechanism – holding fixed the other agents' reports and subject to the average of agent n's reports $\in [\underline{\mu}^n, \overline{\mu}^n]$. In the linked VCG mechanism, the principal pays agent n the difference, $y - x$.

Like the unlinked VCG mechanism, the linked VCG mechanism is an efficient ex-post mechanism. In fact, since $[\underline{\mu}^n, \overline{\mu}^n]$ is path-connected, the envelope theorem implies that the linked VCG mechanism is the cheapest efficient ex-post mechanism. An agent's transfer in the linked VCG mechanism is bounded above by his own maximum contribution to surplus, rather than equalling the sum of everyone else's contributions to surplus as in the unlinked VCG mechanism. This means when an agent is small, in the sense that he contributes at most a small fraction to surplus, the savings from linking are significant. Computing an agent's transfer is straightforward and, when he is small, can be done via a simple geometric algorithm.

The restriction that each agent n's T-date average endowment must lie in $[\underline{\mu}^n, \overline{\mu}^n]$ is meant to be a stylization, where the "unlikely" endowment processes with T-date averages lying outside $[\underline{\mu}^n, \overline{\mu}^n]$ have been removed. The comparison of the linked and unlinked VCG mechanisms shows how such trimming of the underlying type space can greatly improve the solution to a mechanism design problem.

In the last part of the paper, I relax the stylized restriction on average endowments, letting the vector of endowment processes go on indefinitely, potentially taking any value. T, $\{\underline{\mu}^n\}, \{\overline{\mu}^n\}$ are now parameters of the linked VCG mechanism. By varying T, $\{\underline{\mu}^n\}, \{\overline{\mu}^n\}$, the principal controls some common knowledge lower bound, $1-\delta$, on the probability that every agent n's T-date average endowment $\in [\underline{\mu}^n, \overline{\mu}^n]$. I then argue, as δ tends to 0, the probability that all agents have almost ex-post incentives to tell the truth at all dates under the linked VCG mechanism tends to 1. Assuming agents tell the truth whenever it is almost ex-post optimal to do so, the attractive properties of the linked VCG mechanism are approximately preserved when the stylized restriction on the vector of average endowments is relaxed.

As an example, I take a model and scale up its size, with the number of agents and the stocks of the R durable resources increasing proportionally. Assuming agents are patient, I then derive conditions on the endowments' strength of ergodicity – across both agents and dates – such that if it is common knowledge they are satisfied, then in the limit, as the number of agents tends to infinity, an appropriately linked VCG mechanism achieves allocative efficiency at a cost-to-surplus ratio of zero. The conditions are satisfied if, for example, endowments are iid across agents and dates. In contrast, the cost-to-surplus ratio of the unlinked VCG mechanism goes to infinity.

One application of the linked VCG mechanism 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 with the appropriate skills through some dynamic mechanism. See Smet, Lund and Schaninger (2016).

This problem can be viewed through my model: The principal corresponds to the organization's headquarters and the agents correspond to various departments. Projects are departmental tasks while resource types are employee skills. The stocks of durable resources are the organization's employee labor pool parameterized by skill-hours. For example, if an employee's skill is programming, and a date is four weeks, then this employee could represent 4 weeks $\times$ 40 hours/week = 160 hours of the programming resource that can be flexibly allocated, each date, across a variety of projects. Transfers from the principal to agents correspond to incentive pay for department managers.

My work on the linked VCG mechanism over stylized endowment domains is related to the work of Holmström (1979) on VCG mechanisms over restricted preference domains. Holmström's intent was to demonstrate, in a static setting and in the context of efficient dominant-strategy implementation, that the necessity of VCG mechanisms does not rely on agents having "universal" preference domains as in, say, Green and Laffont (1977). In contrast, my stylized restriction that the T-date average of agent n's endowment process must lie in some confidence interval $[\underline{\mu}^n, \overline{\mu}^n]$ is motivated by ergodic considerations. Also related is Bergemann and Välimäki (2010), who study dynamic VCG mechanisms in a Markovian setting.

The restriction on the T-date average introduces a budget that limits how much value an agent can claim over time. 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 construction of a linked VCG mechanism with cost-to-surplus ratio of zero in the many agents case reveals a surprising connection between budget mechanisms and VCG mechanisms.

Jackson and Manelli (1997) show how the attractive properties of the market mechanism are approximately preserved when the price-taking assumption is relaxed. I perform a similar exercise by showing how the attractive properties of the linked VCG mechanism are approximately preserved when the stylized restriction on the vector of endowment processes is relaxed. In particular, I develop a notion of almost ex-post incentives and show that, with high probability, all agents have almost ex-post incentives to tell the truth at all times under the linked VCG mechanism. Recently, Lee (2017) and Azevedo and Budish (2019) have explored related notions of approximate strategy-proofness.

2 A Model of Repeated Resource Allocation

A principal (she) possesses R kinds of durable, divisible resources. Her stock of resource $r = 1, \ldots R$ is $\overline{q}(r) > 0$. Let $\overline{q} := (\overline{q}(r))_{r=1}^R$. A **project** specifies a resource input $\theta = (\theta(r))_{r=1}^R \in [0, \infty)^R - \{0\}$ and a positive scalar payoff. Projects are divisible, so I normalize all project payoffs to 1 and identify a project type by its input.

The domain of project types is a finite set $\Theta = \{\theta_i\}_{i=1,\ldots I}$ satisfying two conditions:

1. $I > R$, and for $i \leq R$, $\theta_i(r) > 0 \Leftrightarrow r = i$;
2. Any subset of Θ lying on a hyperplane with normal vector $p \in (0, \infty)^R$ is linearly independent.

For condition 1, a useful benchmark to keep in mind is when each type θ_i project, $i \leq R$, uses a large amount of resource i, and should therefore be thought of as an inefficient project, to be invested in as a last resort. The second condition is generic. Neither condition is crucial, but they do simply the analysis. If condition 1 were not satisfied, an efficient use of resources would typically leave some resources unused, even if there were a plentiful supply of each project type. If condition 2 were not satisfied, there could be multiple different ways to efficiently use resources.

By combining conditions 1 and 2 with an assumption, shortly, about projects being sufficiently plentiful, I ensure that there is a unique way to efficiently use resources, and under this efficient usage, all resources are exhausted. See Figure 1.

Fix a time horizon $T \geq 1$. At each date $t = 1, 2, \ldots T$, the principal is endowed with an infinite amount of the type θ_i project for $i \leq R$ and none of the other projects. Let μ_t^0 denote the principal's date t endowment.

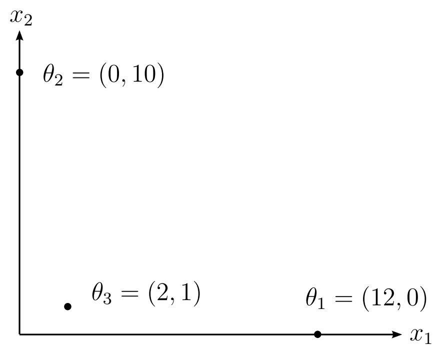

(a) $R = 2$, $I = 3$.

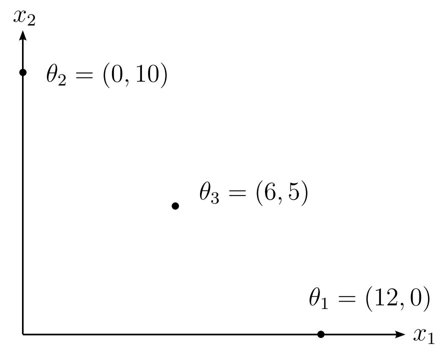

(b) $R = 2$, $I = 3$, $\overline{q} = (3, 2.5)$.

Figure 1: In 1a, suppose there is an infinite supply of θ_1 and θ_2 projects and 1 unit of θ_3 projects. To maximize payoff, one will invest in θ_3 projects until all such projects are exhausted or one of the resources is exhausted. At that point, any left over resources are then invested in θ_1 or θ_2 projects.

In 1b, condition 2 is violated. Suppose there is an infinite supply of θ_1 and θ_2 projects and 1 unit of θ_3 projects. There are multiple ways to maximize payoff. One could invest in 0.5 units of θ_3 projects, or invest in 0.25 units each of θ_1 and θ_2 projects.

There are $N \geq 1$ agents (he). At each date t, each agent $n = 1, \ldots N$ privately observes his date t endowment of projects, $\mu_t^n := (\mu_{ti}^n)_{i=1}^I$. Let $\Omega := [0, \infty)^I$ denote the set of possible agent endowments, with Borel sigma algebra $\mathcal{F}$. Agent n's endowment process is the T-row vector

$$\boldsymbol{\mu}^n = \begin{bmatrix} \mu_1^n & \mu_2^n & \ldots & \mu_T^n \end{bmatrix}.$$

Notation: A bold letter with a superscript (e.g. $\boldsymbol{\mu}^n$) denotes a time-related T-row vector associated with an agent; that bolded letter, by itself (e.g. $\boldsymbol{\mu}$), denotes the $N \times T$ matrix that collects the row vectors across all agents; that bolded letter, with a subscript (e.g. $\boldsymbol{\mu}_t$), denotes an N-column vector of the matrix. Appending $|_t$ to a matrix denotes a new matrix consisting of the first t columns of the original matrix.

Exceptions: When an object is indexed by agent but not time, I use a bolded letter, without a subscript, to denote the N-column vector. $\mathbf{1}$ denotes either a row or column vector with 1 in each entry and length depending on context.

A (direct) mechanism $(\boldsymbol{A}, \boldsymbol{W})$ consists of

- A message space, Ω, for each agent n and date t;
- An allocation $A_t^n : \boldsymbol{\Omega}|_t \to [0, \infty)^R$ for each agent n and date t;
  - $\boldsymbol{A}_t$ satisfies the feasibility constraint $\mathbf{1}\boldsymbol{A}_t \leq \overline{q}$;
  - Define $A_t^0 := \overline{q} - \mathbf{1}\boldsymbol{A}_t$ to be the leftover resources allocated to the principal;

- A transfer $W^n : \mathbf{\Omega} \to [0, \infty)$ for each agent n.

A strategy $\boldsymbol{\sigma}^n$ for agent n consists of mappings $\sigma_t^n : \mathbf{\Omega}^n|_t \times \mathbf{\Omega}^{-n}|_{t-1} \to \Omega$ for each date t, where $\mathbf{\Omega}^{-n}|_0 := \{\emptyset\}$. $\boldsymbol{\sigma}$ can be viewed as a map from $\mathbf{\Omega}$ to itself: Define $\boldsymbol{\sigma}(\boldsymbol{\mu})_1^n := \sigma_1^n(\mu_1^n \times \emptyset)$ for all n, and, for $t > 1$, recursively define $\boldsymbol{\sigma}(\boldsymbol{\mu})_t^n = \sigma_t^n(\boldsymbol{\mu}^n|_t \times \boldsymbol{\sigma}(\boldsymbol{\mu})^{-n}|_{t-1})$.

Given $q \in (0, \infty)^R$ resources and $\mu \in [0, \infty]^R \times [0, \infty)^{R-I}$ projects, let

$$v(q, \mu) := \max_{\alpha \in [0,\mu],\ \sum_i \alpha_i \theta_i \leq q} \sum_i \alpha_i. \tag{1}$$

Given $(\boldsymbol{A}, \boldsymbol{W})$ and $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, player n's contribution to surplus, the surplus, the cost to the principal, and the payoff to agent n are, respectively,

$$\begin{aligned}
S^n(\boldsymbol{\sigma}(\boldsymbol{\mu})) &:= \sum_{t=1}^{T} \beta^{t-1} v\left(A_t^n(\boldsymbol{\sigma}(\boldsymbol{\mu})|_t), \mu_t^n\right) \\
S(\boldsymbol{\sigma}(\boldsymbol{\mu})) &:= \sum_{n=0}^{N} S^n(\boldsymbol{\sigma}(\boldsymbol{\mu})), \\
C(\boldsymbol{\sigma}(\boldsymbol{\mu})) &:= \beta^{T-1} \mathbf{1} \boldsymbol{W}(\boldsymbol{\sigma}(\boldsymbol{\mu})), \\
U^n(\boldsymbol{\sigma}(\boldsymbol{\mu})) &:= S^n(\boldsymbol{\sigma}(\boldsymbol{\mu})) + \beta^{T-1} W^n(\boldsymbol{\sigma}(\boldsymbol{\mu})).
\end{aligned}$$

Fix $\underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}} \in \Omega^N$ with $\underline{\mu}^n \leq \overline{\mu}^n$ for all n. Define

$$\mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) = \left\{ \boldsymbol{\mu} \in \mathbf{\Omega} \;\middle|\; \underline{\mu}^n \leq \frac{\sum_{t=1}^{T} \mu_t^n}{T} \leq \overline{\mu}^n \ \forall n \right\}.$$

I assume it is common knowledge that $\boldsymbol{\mu}$ must lie in $\mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$. One should think of this assumption as a stylization where any $\boldsymbol{\mu} \notin \mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ is considered "unlikely" and has been removed from the type space.

Definition. *Given* $(\boldsymbol{A}, \boldsymbol{W})$, $\boldsymbol{\sigma}$ *is an ex-post equilibrium if, for all* $\boldsymbol{\mu} \in \mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, n, *and* $\hat{\boldsymbol{\sigma}}^n$, $U^n(\boldsymbol{\sigma}(\boldsymbol{\mu})) \geq U^n((\hat{\boldsymbol{\sigma}}^n, \boldsymbol{\sigma}^{-n})(\boldsymbol{\mu}))$.

Given $\boldsymbol{\sigma}^n$, whenever $\boldsymbol{\mu}$ satisfies $\sigma_t^n(\boldsymbol{\mu}^n|_t \times \mathbf{\Omega}^{-n}|_{t-1}) = \mu_t^n$ for all t, agent n is said to be employing the truth-telling strategy. The truth-telling strategy profile consists of each agent n employing the truth-telling strategy for all $\boldsymbol{\mu} \in \mathbf{\Omega}$.

Without loss of generality, I restrict attention to mechanisms where the truth-telling strategy profile is an ex-post equilibrium and I call them ex-post mechanisms. $\boldsymbol{A}$ is **efficient** if, for all $\boldsymbol{\mu} \in \mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, $S(\boldsymbol{\mu}) = \sum_{t=1}^{T} \beta^{t-1} v(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0)$. An ex-post mechanism $(\boldsymbol{A}, \boldsymbol{W})$ is efficient if $\boldsymbol{A}$ is efficient. Given efficient ex-post mechanisms, $(\boldsymbol{A}, \boldsymbol{W})$ and $(\hat{\boldsymbol{A}}, \hat{\boldsymbol{W}})$, $(\boldsymbol{A}, \boldsymbol{W})$ is weakly cheaper than $(\hat{\boldsymbol{A}}, \hat{\boldsymbol{W}})$ if, for all $\boldsymbol{\mu} \in \mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, $C(\boldsymbol{\mu}) \leq \hat{C}(\boldsymbol{\mu})$. The mechanism design problem is to find a cheapest efficient ex-post mechanism.

3 Divisible Multi-Dimensional Knapsack Problem

The solution to the mechanism design problem begins with the solution to (1). This is the divisible multi-dimensional knapsack problem, where q parameterizes the multi-dimensional knapsack and μ represents the set of divisible objects.

Given q and μ, call any $\alpha \in [0, \mu]$ satisfying $\sum_i \alpha_i \theta_i \leq q$ an *investment*. α is *efficient* if it solves (1). α is *exhaustive* if $\sum_i \alpha_i \theta_i = q$. α is *linear* if there exists a $p \in (0, \infty)^R$ such that $p \cdot \theta_i < 1 \Rightarrow \alpha_i = \mu_i$ and $p \cdot \theta_i > 1 \Rightarrow \alpha_i = 0$, in which case, call p a price associated with α. Given price p, a type i project can be thought of as having cost $p \cdot \theta_i$. In a linear investment, a project is invested in (not invested in) if its payoff is strictly greater than (strictly less than) its cost under an associated price. Note, there are no restrictions on investments in projects whose cost equals payoff.

Lemma 1. *Given q and $\mu \geq \mu_t^0$, there exists a unique investment, $\alpha(q, \mu)$, that is both exhaustive and linear. $\alpha(q, \mu)$ is the unique efficient investment.*

Proof. See appendix. □

To gain some intuition for this result, consider an arrangement of Θ depicted in Figure 2a. Suppose there are two distinct exhaustive and linear investments, α^* and $\hat{\alpha}$, where α^* ($\hat{\alpha}$) involves investing in (not investing in) any project strictly under (strictly above) the dotted line labelled α^* ($\hat{\alpha}$).

Let X denote the intersection of the two dotted lines. Let us assume, for simplicity, that no project type is located in X. The two dotted lines create two triangles in which are located $\{\theta_i, \theta_j, \theta_k, \theta_l, \theta_m, \theta_n\}$. By definition,

- $\alpha_i^* - \hat{\alpha}_i \geq 0$, $\alpha_j^* - \hat{\alpha}_j \geq 0$, $\alpha_k^* - \hat{\alpha}_k \geq 0$,
- $\alpha_l^* - \hat{\alpha}_l \leq 0$, $\alpha_m^* - \hat{\alpha}_m \leq 0$, $\alpha_n^* - \hat{\alpha}_n \leq 0$,
- $\alpha_h^* - \hat{\alpha}_h = 0$ for all $h \in \{1, 2, \ldots I\} - \{i, j, k, l, m, n\}$.

Moreover, since α^* and $\hat{\alpha}$ are distinct, at least one of the differences in the first two bullet points must be nonzero.

If X is excluded, then the two triangles are two disjoint convex sets, and, therefore, can be separated by a hyperplane. It is easy to see such a separating hyperplane can be chosen so that it goes through the origin. See Figure 2b. Let z be a normal vector on the side of, say, the ijk triangle. By construction, $z \cdot \theta_i, z \cdot \theta_j, z \cdot \theta_k > 0 > z \cdot \theta_l, z \cdot \theta_m, z \cdot \theta_n$. Thus, $z \cdot (\alpha_h^* - \hat{\alpha}_h)\theta_h \geq 0$ for all $h \in \{1, 2, \ldots I\}$, and at least one of them is positive for $h \in \{i, j, k, l, m, n\}$. And now, by the fact that both α^* and $\hat{\alpha}$ are exhaustive, we have a contradiction because $0 = z \cdot (q - q) = z \cdot \left[\sum_{h=1}^I \alpha_h^* \theta_h - \sum_{h=1}^I \hat{\alpha}_h \theta_h\right] = \sum_{h=1}^I z \cdot (\alpha_h^* - \hat{\alpha}_h)\theta_h > 0$.

To show existence of an exhaustive and linear investment, it suffices to show that an efficient investment must be exhaustive and linear. That it must be exhaustive is

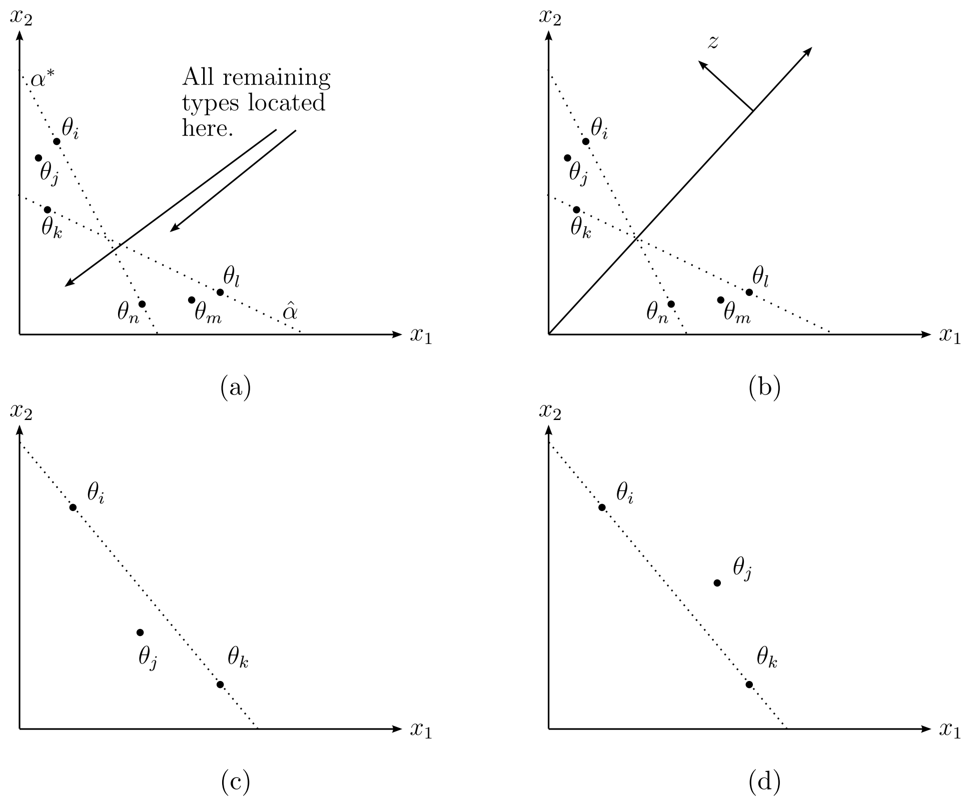

Figure 2

obvious. So let us consider linearity. Note, in a linear investment, the set of project types that are partially invested in must lie on a hyperplane. Suppose this were not the case. Figures 2c and 2d depict two such non-linearities, with type $\theta_i, \theta_j, \theta_k$ projects all partially invested in.

In 2c, let $\lambda \in [0,1]$ satisfy $\lambda\theta_i + (1-\lambda)\theta_k > \theta_j$. Now consider the following slight modification to the investment: Increase the investment in θ_j projects slightly by dx. Simultaneously, decrease the investment in θ_i and θ_k projects by $\lambda dx\theta_i$ and $(1-\lambda)dx\theta_k$, respectively. The net change in payoff is $dx - \lambda dx - (1-\lambda)dx = 0$, while the net change in resource use is $dx\theta_j - \lambda dx\theta_i - (1-\lambda)dx\theta_k < 0$. Now there is a little bit of each resource unused. Invest them in some more θ_j projects, leading to a higher payoff. A similar argument can be used to rule out 2d.

Let $QM := \{(q,\mu) \mid \mu \geq \mu_t^0\}$. By identifying QM with $(0,\infty)^R \times [0,\infty)^{I-R}$, we can view QM as a cone equipped with the norm $\|\cdot\|_\infty$. Given Lemma 1, for any $(q,\mu) \in QM$, let $p(q,\mu)$ denote the set of prices associated with efficient investment $\alpha(q,\mu)$.

Corollary 1. $p(\cdot,\cdot)$ *is a homogenous correspondence of degree 0.*

Proof. Let $\lambda > 0$. That $\alpha(q,\mu)$ is linear with respect to μ and exhausts q implies $\lambda\alpha(q,\mu)$ is linear with respect to $\lambda\mu$ and exhausts λq. Thus, $\alpha(\lambda q, \lambda\mu) = \lambda\alpha(q,\mu)$.

This implies $p \in p(\lambda q, \lambda\mu) \Leftrightarrow p \in p(q, \mu)$. □

Lemma 2. *$p(\cdot,\cdot)$ is generically unique and locally constant: $\{(q,\mu) \in QM \mid |p(q,\mu)| = 1\}$ is an open dense subset of $(QM, \|\cdot\|_\infty)$. If $|p(q,\mu)| = 1$, then there exists a $d > 0$ such that for all $(\tilde{q},\tilde{\mu}) \in QM$ satisfying $\|(\tilde{q},\tilde{\mu}) - (q,\mu)\|_\infty < d$, $p(\tilde{q},\tilde{\mu}) = p(q,\mu)$.*

Proof. See appendix. □

The following lemma bounds the second partial of $v(\cdot,\cdot)$ and will allow us to use an envelope theorem later on.

Lemma 3. *$v(q,\cdot)$ is concave, and, for all $\mu, \nu \geq \mu_t^0$ and $\lambda^* \in (0,1]$, $\frac{d}{d\lambda}^- \big|_{\lambda=\lambda^*} v(q, (1-\lambda)\mu + \lambda\nu) < \sum_{i \leq R} \frac{\overline{q}(i)}{\theta_i(i)} + \sum_{i>R} \mu_i \vee \nu_i$.*

Proof. See appendix. □

4 Linking VCG Mechanisms

Given $\boldsymbol{\mu}$, Lemma 1 implies the efficient investment $\alpha(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0)$ each date is unique. Nevertheless, the efficient allocation $\boldsymbol{A}_t$ each date is usually not unique. This is because whenever $\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) \in \left(0, \sum_{n \geq 0} \mu_{ti}^n\right)$, there is some flexibility in deciding what portion of resources meant for type θ_i projects go to each player.

Definition. *Given $\boldsymbol{\mu}_t$ and i, define $N(\boldsymbol{\mu}_t, i) \in \{0, \dots N\}$ to be the index satisfying $\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) \in (\sum_{m=N(\boldsymbol{\mu}_t,i)+1}^N \mu_{ti}^m, \sum_{m=N(\boldsymbol{\mu}_t,i)}^N \mu_{ti}^m]$. If $\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) = 0$, then set $N(\boldsymbol{\mu}_t, i) = N$. For $n = 0, \dots N$,*

$$_{ord}A_t^n(\boldsymbol{\mu}|_t) := \sum_{i=1}^{I} \left[1_{n > N(\mathbf{1}\boldsymbol{\mu}_t, i)} \mu_{ti}^n + 1_{n = N(\mathbf{1}\boldsymbol{\mu}_t, i)} \left(\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) - \sum_{m=n+1}^{N} \mu_{ti}^m \right) \right] \theta_i.$$

$_{ord}\boldsymbol{A}$ is the efficient allocation that serves everyone in order according to their index, with agent N being served first and the principal being served last. From now on, I restrict attention to the efficient allocation $_{ord}\boldsymbol{A}$.

Definition. *The unlinked VCG mechanism $(_u\boldsymbol{A}, {_u\boldsymbol{W}})$ is defined as follows: For each agent n,*

1. $_uA_t^n(\boldsymbol{\mu}|_t) = {_{ord}A_t^n}(\boldsymbol{\mu}|_t)$,
2. $_uW^n(\boldsymbol{\mu}) = \sum_{t=1}^T \sum_{m \neq n,\ m=0}^N \beta^{t-T} v(_uA_t^m(\boldsymbol{\mu}|_t), \mu_t^m)$.

Proposition 1. *$(_u\boldsymbol{A}, {_u\boldsymbol{W}})$ is an efficient ex-post mechanism.*

Proof. The classic result that VCG mechanisms are efficient and strategy-proof when $T = 1$ implies that $(_u\boldsymbol{A}, {_u\boldsymbol{W}})$ is an efficient ex-post mechanism for all T. □

For $T > 1$, $({}_u\boldsymbol{A}, {}_u\boldsymbol{W})$ need not be strategy-proof. Fix an n, and consider the $\boldsymbol{\sigma}^{-n}$ where at all dates $t > 1$, all agents $m \neq n$ copy agent n's date 1 report. At date 1, agent n is better off reporting a sufficiently huge endowment rather than the truth.

For each $\boldsymbol{\mu} \in \boldsymbol{\Omega}$, define $[\boldsymbol{\mu}] \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ as follows: For each n, i, and t,

$$[\mu]^n_{ti} := \begin{cases} \mu^n_{ti} & \text{if } (\boldsymbol{\mu}^n|_t\mathbf{1})_i \leq T\overline{\mu}^n_i \text{ and } t < T \\ T\overline{\mu}^n_i - (\boldsymbol{\mu}^n|_{t-1}\mathbf{1})_i & \text{if } (\boldsymbol{\mu}^n|_{t-1}\mathbf{1})_i \leq T\overline{\mu}^n_i < (\boldsymbol{\mu}^n|_t\mathbf{1})_i \\ 0 & \text{if } T\overline{\mu}^n_i < (\boldsymbol{\mu}^n|_{t-1}\mathbf{1})_i. \end{cases}$$

and

$$[\mu]^n_{Ti} := \max\{\mu^n_{Ti}, T\underline{\mu}^n_{Ti} - (\boldsymbol{\mu}^n|_{T-1}\mathbf{1})_i\} \;\text{ if } (\boldsymbol{\mu}^n|_T\mathbf{1})_i \leq T\overline{\mu}^n_i.$$

$[\cdot]$ is an $\mathcal{F}_t$-measurable retraction from $\boldsymbol{\Omega}$ to $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$.

Definition. *The linked VCG mechanism $({}_l\boldsymbol{A}, {}_l\boldsymbol{W})$ is defined as follows: For each agent n,*

1. ${}_lA^n_t(\boldsymbol{\mu}|_t) = {}_uA^n_t([\boldsymbol{\mu}]|_t)$,
2. ${}_lW^n(\boldsymbol{\mu}) = {}_uW^n([\boldsymbol{\mu}]) - \min_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})} {}_uW^n([\boldsymbol{\mu}]^{-n}, \hat{\boldsymbol{\mu}}^n)$.

Theorem 1. *$({}_l\boldsymbol{A}, {}_l\boldsymbol{W})$ is an efficient ex-post mechanism. It is weakly cheaper than every efficient ex-post $(\boldsymbol{A}, \boldsymbol{W})$ satisfying $\boldsymbol{A}(\boldsymbol{\mu}) = {}_u\boldsymbol{A}(\boldsymbol{\mu})$ for all $\boldsymbol{\mu} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$.*

Lemma 4. *If an efficient ex-post $(\boldsymbol{A}, \boldsymbol{W})$ satisfies $\boldsymbol{A}(\boldsymbol{\mu}) = {}_u\boldsymbol{A}(\boldsymbol{\mu})$ for all $\boldsymbol{\mu} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, then, for all n, $W^n - {}_uW^n$ depends on $\boldsymbol{\mu}$ only up to $\boldsymbol{\mu}^{-n}$ on $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$.*

Proof. Fix an efficient ex-post $(\boldsymbol{A}, \boldsymbol{W})$ satisfying $\boldsymbol{A}(\boldsymbol{\mu}) = {}_u\boldsymbol{A}(\boldsymbol{\mu})$ for all $\boldsymbol{\mu} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, an agent n, and a pair $\boldsymbol{\mu}, \tilde{\boldsymbol{\mu}} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ with $\boldsymbol{\mu}^{-n} = \tilde{\boldsymbol{\mu}}^{-n}$. Define $V^n : \boldsymbol{\Omega}^n \times [0, 1] \to [0, \infty)$ by

$$V^n(\hat{\boldsymbol{\mu}}^n, \lambda) = \sum_{t=1}^{T} \beta^{t-1} v\left(A^n_t(\boldsymbol{\mu}^{-n}_t, \hat{\mu}^n_t),\; \lambda\tilde{\mu}^n_t + (1-\lambda)\mu^n_t\right) + \beta^{T-1} W^n(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n)$$

Lemma 3 implies $V^n(\hat{\boldsymbol{\mu}}^n, \lambda)$ is absolutely continuous with respect to λ with a partial derivative that is uniformly bounded above. Thus, by Theorem 2 of Milgrom and Segal (2002),

$$\begin{aligned} V^n(\tilde{\boldsymbol{\mu}}^n, 1) - V^n(\boldsymbol{\mu}^n, 0) &= \int_0^1 V^n_2(\lambda\tilde{\boldsymbol{\mu}}^n + (1-\lambda)\boldsymbol{\mu}^n, \lambda) d\lambda \\ &= \int_0^1 {}_uV^n_2(\lambda\tilde{\boldsymbol{\mu}}^n + (1-\lambda)\boldsymbol{\mu}^n, \lambda) d\lambda \\ &= {}_uV^n(\tilde{\boldsymbol{\mu}}^n, 1) - {}_uV^n(\boldsymbol{\mu}^n, 0). \end{aligned}$$

This then implies,

$$\begin{aligned}\beta^{T-1}(W^n(\tilde{\boldsymbol{\mu}}) - {}_uW^n(\tilde{\boldsymbol{\mu}})) &= V^n(\tilde{\boldsymbol{\mu}}^n, 1) - {}_uV^n(\tilde{\boldsymbol{\mu}}^n, 1) \\ &= V^n(\boldsymbol{\mu}^n, 0) - {}_uV^n(\boldsymbol{\mu}^n, 0) \\ &= \beta^{T-1}(W^n(\boldsymbol{\mu}) - {}_uW^n(\boldsymbol{\mu}))\end{aligned}$$

□

Proof of Theorem 1. The first part of Theorem 1 follows from Proposition 1. To prove the second part, it suffices to show ${}_lW^n(\boldsymbol{\mu}) \leq W^n(\boldsymbol{\mu})$ for all n and $\boldsymbol{\mu} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$.

Suppose not. Then there exists n and $\boldsymbol{\mu} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ satisfying $W^n(\boldsymbol{\mu}) = {}_lW^n(\boldsymbol{\mu}) - \delta$ for some $\delta > 0$. By definition of ${}_lW^n$, there exists a $\hat{\boldsymbol{\mu}} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ satisfying $\hat{\boldsymbol{\mu}}^{-n} = \boldsymbol{\mu}^{-n}$ and ${}_lW^n(\hat{\boldsymbol{\mu}}) = 0$. Lemma 4 implies

$$\begin{aligned}W^n(\hat{\boldsymbol{\mu}}) &= W^n(\hat{\boldsymbol{\mu}}) - {}_uW^n(\hat{\boldsymbol{\mu}}) + \min_{\tilde{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})} {}_uW^n(\hat{\boldsymbol{\mu}}^{-n}, \tilde{\boldsymbol{\mu}}^n) \\ &= W^n(\boldsymbol{\mu}) - {}_uW^n(\boldsymbol{\mu}) + \min_{\tilde{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})} {}_uW^n(\boldsymbol{\mu}^{-n}, \tilde{\boldsymbol{\mu}}^n) \\ &= W^n(\boldsymbol{\mu}) - {}_lW^n(\boldsymbol{\mu}) = -\delta < 0.\end{aligned}$$

Contradiction. □

Proposition 1 and Theorem 1 are in fact true no matter what efficient allocation ${}_u\boldsymbol{A}$ is set to be. The choice of ${}_{ord}\boldsymbol{A}$ is made mostly for some computational benefits later on. Note, changing the efficient allocation of the linked VCG mechanism changes the cost function, $C(\cdot)$. For example, consider the efficient allocation, ${}_{p-ord}\boldsymbol{A}$, that serves the agents in order like ${}_{ord}\boldsymbol{A}$ but serves the principal first. The linked VCG mechanism with ${}_{ord}\boldsymbol{A}$ is weakly cheaper than the linked VCG mechanism with ${}_{p-ord}\boldsymbol{A}$. In general, however, linked VCG mechanisms with different efficient allocations cannot be completely ordered based on cheapness.

Theorem 1 implies that linking VCG mechanisms leads to cost savings for the principal. How much cost savings depends on the parameters. For example, if $\overline{\boldsymbol{\mu}} \gg \underline{\boldsymbol{\mu}}$, then the cost savings are zero.

For now, I just want to highlight a *qualitative difference* in the size of transfers between the linked VCG mechanism and the unlinked VCG mechanism. For simplicity, let us specialize to $R = T = 1$. In addition, let us depart slightly from the model by assuming that Θ is equal to $(0, \infty)$ rather than a finite subset of $(0, \infty)$, and that endowments are densities on $(0, \infty)$ – this departure makes drawing pictures easier. The efficient allocation sets a threshold $\overline{\theta}$ and invests in all projects with type $\theta \leq \overline{\theta}$. $\overline{\theta}$ is set so that all of the resource is exhausted.

Figures 3a and 3b compare the efficient allocations given two agent n endowments $\boldsymbol{\mu}^n, \hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(1, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, holding fixed all other agents' endowments at some $\boldsymbol{\mu}^{-n} \in \boldsymbol{\Omega}^{-n}(1, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$. As depicted, $\hat{\boldsymbol{\mu}}^n$ is a larger endowment than $\boldsymbol{\mu}^n$, pushing the threshold $\overline{\theta}$

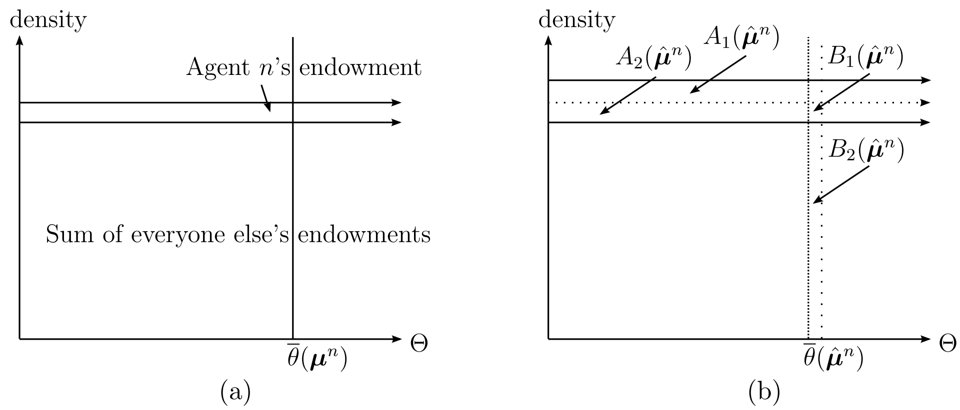

Figure 3

to the left. Since projects further to the left use fewer resources, exhaustiveness implies $A_1(\hat{\boldsymbol{\mu}}^n)$ has greater area than $B_1(\hat{\boldsymbol{\mu}}^n) \cup B_2(\hat{\boldsymbol{\mu}}^n)$. Now consider the difference ${}_uW^n(\boldsymbol{\mu}) - {}_uW^n(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n)$. It is the area of $B_2(\hat{\boldsymbol{\mu}}^n)$, which, given the previous observation, is less than the area of $A_1(\hat{\boldsymbol{\mu}}^n) \cup A_2(\hat{\boldsymbol{\mu}}^n)$, which, by definition, is agent n's contribution to surplus given $\hat{\boldsymbol{\mu}}^n$. Since ${}_lW(\boldsymbol{\mu})$ is the supremum of this difference across all $\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(1, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, agent n's transfer under the linked VCG mechanism is bounded above by his own maximum contribution to surplus, rather than equalling the sum of everyone else's contributions to surplus, as is the case under the unlinked VCG mechanism.

Thus, if an agent is small, in the sense that he contributes at most a small fraction to surplus, then the cost savings of linking are significant. See Section 6 for details.

5 From Ex-Post to ε-Ex-Post Incentives

So far, we have assumed it is common knowledge the T-date average endowment vector must lie in some interval $[\underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}]$. Such a model should be thought of as a stylization of some "true" model where the vector of endowment processes can go on indefinitely, potentially taking any value, but $(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ have been chosen in a way so that it is common knowledge that the T-date average almost certainly lies in $[\underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}]$.

I now construct such a "true" model: Change the stylized model so that each agent n now has an endowment process that goes on indefinitely. Redefine $(\boldsymbol{\Omega}, \boldsymbol{\mathcal{F}})$ accordingly and let $\mathbf{P}$ denote the true governing probability. $\mathbf{P}$ *is not necessarily common knowledge.* However, I assume it is common knowledge each agent n knows his own marginal, P^n.

T, $\underline{\boldsymbol{\mu}}$, and $\overline{\boldsymbol{\mu}}$ are no longer parameters of the model. The principal chooses T as part of choosing the mechanism $(\boldsymbol{A}, \boldsymbol{W})$. In particular, linked VCG mechanisms are parameterized by a triple $(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ satisfying $\underline{\boldsymbol{\mu}} \leq \overline{\boldsymbol{\mu}}$. Finally, fix an $\varepsilon > 0$ interpreted to be small. This completes the description of the true model.

Fix an arbitrary linked VCG mechanism $(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$. Let $\delta \in [0, 1]$ be a number such that it is common knowledge $\mathbf{P}(\boldsymbol{\Omega}(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})) \geq 1-\delta$. Such a δ always exists – for example, $\delta = 1$. Depending on how much the players know about $\mathbf{P}$ and how $(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})$ is chosen, lower δ may also exist. For example, if $\mathbf{P}$ is common knowledge and $(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})$ has been chosen in a way so that $\mathbf{P}(\boldsymbol{\Omega}(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})) = .95$, then δ can be set to anything ≥ 0.05. Here, we will simply take δ as given.

What can the principal expect from using the linked VCG mechanism $(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})$, particularly when δ is small? Were we in the stylized model $(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})$, Theorem 1 would tell us that the principal can expect all agents to tell the truth, the allocation to be efficient, and the cost to be weakly lower than that of any other efficient ex-post mechanism. But here in the true model, how far off are we from truth-telling being an ex-post equilibrium? Can the principal expect the allocation to be at least almost efficient? How different is the cost of using the mechanism?

Let $\overline{v}$ denote the finite value $v(\overline{q},\infty\mathbf{1})$. For now, let us assume $N \geq 2$.

Definition. *Given a $c^n \in [0,1]$, an agent n strategy $\boldsymbol{\sigma}^n$ is said to be reasonable if $\sigma^n_t(\boldsymbol{\mu}^n|_t \times \boldsymbol{\Omega}^{-n}|_{t-1}) = \mu^n_t$ whenever*

$$\left[P^n\left(\boldsymbol{\mu}^n \notin \boldsymbol{\Omega}^n(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}}) \;\middle|\; \boldsymbol{\mu}^n|_s\right) + c^n\right]\beta^{1-s}T\overline{v} < \varepsilon \tag{2}$$

for all $s \leq t$.

In words, an agent strategy is reasonable if it is truth-telling until (2) is violated.

The c^n in the definition is meant to be interpreted as a conjecture: Agent n conjectures the probability that all other agents employ the truth-telling strategy is at least $1-c^n$. Given a profile of conjectures, $\boldsymbol{c}$, a strategy profile $\boldsymbol{\sigma}$ is reasonable if each $\boldsymbol{\sigma}^n$ is reasonable given c^n.

To answer the questions posed above, I imagine, before date 1, the principal publicly announces a profile of conjectures, $\boldsymbol{c}$, that is "believable" given δ – where believable is yet to be defined. I posit each agent n will believe c^n and play a reasonable strategy given c^n. I then see what I can deduce about the performance of the linked VCG mechanism $(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})$ based only on the knowledge that the strategy profile being played is reasonable given $\boldsymbol{c}$. The principal and any agent $m \neq n$ need not know if some $\boldsymbol{\sigma}^n$ is reasonable given c^n since they need not know P^n. Thus, whether or not some strategy profile $\boldsymbol{\sigma}$ is reasonable given $\boldsymbol{c}$ need not be known to anyone.

For now, let us take as given that agent n will believe conjecture c^n. Why do I then assume that he will play a reasonable strategy given c^n?

Justification. Imagine agent n enters date 1 with a coarse assessment (P^n, c^n) of the game. I use the word coarse because in a full assessment, the system of beliefs would imply P^n and, together with the strategy profile, would imply the probability that all agents $m \neq n$ employ the truth-telling strategy.

After observing μ^n_1, agent n Bayesian updates P^n to $P^n\big|\boldsymbol{\mu}^n|_1$, and his current coarse assessment becomes $(P^n\big|\boldsymbol{\mu}^n|_1, c^n)$. He could try to further improve this coarse assessment – potentially all the way up to a full assessment. However, any further improvement requires thinking about other agents – e.g. what they know, what they plan to do – beyond just the conjecture c^n. I assume such thinking is at least slightly costly, causing agent n to suffer a disutility $\geq \varepsilon$.

In the event where all other agents employ the truth-telling strategy and $\boldsymbol{\mu}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, telling the truth at all dates is ex-post optimal for agent n. When agent n is deciding whether or not to improve his current coarse assessment $(P^n \big| \boldsymbol{\mu}^n|_1, c^n)$, he believes an upper bound on the probability that that event does not occur is

$$P^n\left(\boldsymbol{\mu}^n \notin \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \;\middle|\; \boldsymbol{\mu}^n|_1\right) + c^n.$$

Now suppose (2) is satisfied with $s = 1$. Since ${}_lU^n(\boldsymbol{\sigma}(\boldsymbol{\mu})) \leq {}_uU^n(\boldsymbol{\sigma}(\boldsymbol{\mu})) \in [0, T\overline{v}]$ for all $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, agent n believes employing the truth-telling strategy is within ε of the payoff he would receive if he could make the ex-post optimal report. Since ε is small, it stands to reason he will not further improve his current coarse assessment at date 1, and will simply tell the truth at date 1.

The agent then enters date 2 with coarse assessment $(P^n \big| \boldsymbol{\mu}^n|_1, c^n)$ and the assumption that agent n plays a reasonable strategy given c^n is now justified by induction. $\square$

Let $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, P^n, c^n)$ denote the event in which agent n is employing the truth-telling strategy under any reasonable strategy given c^n (i.e. the set of $\boldsymbol{\mu}$ such that $\boldsymbol{\mu}^n$ satisfies (2) for all s). Let $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \mathbf{P}, \boldsymbol{c}) := \cap_n \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, P^n, c^n)$ denote the event in which all agents are employing the truth-telling strategy under any reasonable strategy profile given $\boldsymbol{c}$. Given the justification, $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, P^n, c^n)$ can be interpreted as an event in which agent n has ε-ex-post incentives to tell the truth if he believes conjecture c^n. $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \mathbf{P}, \boldsymbol{c})$ can then interpreted as an event in which every agent has ε-ex-post incentives to tell the truth if each agent n believes conjecture c^n.

When $\mathbf{P}(\boldsymbol{\Omega} - \cap_{m \neq n} \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, P^m, c^m)) \leq c^n$, agent n is correct to believe conjecture c^n if each agent $m \neq n$ plays a reasonable strategy given c^m. If this inequality is satisfied for every n, then one could say $\boldsymbol{c}$ is a mutually correct profile of conjectures.

Definition. *$\boldsymbol{c}$ is believable if it is common knowledge that $\boldsymbol{c}$ is mutually correct.*

The trivial profile of conjectures $\mathbf{1}$ is always believable. However, since the condition for mutual correctness utilizes $\{P^n\}_{n=1}^N$, which need not be common knowledge, it is, a priori, not obvious there exist any other believable $\boldsymbol{c}$.

Definition. *For any $z \in (0,1)$, define $\delta(z) = \frac{z(1-z)\varepsilon^2}{(\beta^{1-T}T\overline{v})^2 N} \wedge 1$ and $c(z) = \frac{z\varepsilon}{\beta^{1-T}T\overline{v}} \wedge 1$.*

Notice, the functions $\delta(\cdot)$ and $c(\cdot)$ are defined independently of $\mathbf{P}$.

Proposition 2. *Let $\hat{\mathbf{P}}$ be any probability on $(\boldsymbol{\Omega}, \boldsymbol{\mathcal{F}})$ satisfying $\hat{\mathbf{P}}(\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})) \geq 1 - \delta(z)$. Then $\hat{\mathbf{P}}(\boldsymbol{\Omega} - \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \hat{\mathbf{P}}, c(z)\mathbf{1})) \leq c(z)$.*

Thus, $c(z)\mathbf{1}$ is believable for any z satisfying $\delta(z) \geq \delta$.

If $c(z)\mathbf{1}$ is believable – so that we may assume agents play a reasonable strategy profile given $c(z)\mathbf{1}$ – then in the event, $\mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \mathbf{P}, c(z)\mathbf{1})$, all agents tell the truth. Since this event is a subset of $\mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, if all agents tell the truth in this event, then, in this event, the allocation is efficient and the cost of the mechanism is exactly the same as in the stylized model.

Also, notice, for any $\bar{c} > 0$, there exists a $\overline{\delta} > 0$ such that, for any z satisfying $\delta(z) < \overline{\delta}$, we have $c(z) < \bar{c}$.

Thus, Proposition 2 can be interpreted as saying: When it is common knowledge that the stylized model $(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ is a good approximation of the true model, the principal can expect that, under the linked VCG mechanism $(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, it is almost certain all agents have ε-ex-post incentives to tell the truth, the allocation is efficient and the cost is the same as in the stylized model $(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$. Proposition 2 is the formal answer to the previously posed questions concerning to what extent the attractive properties of the linked VCG mechanism are preserved moving from the stylized model to the true model.

The proof of Proposition 2 makes use of the following result.

Lemma 5. *Let $X = \{X_1, X_2 \dots X_T\}$ be a nonnegative martingale under some probability Q. If X_0 is an upper bound on its expectation, then for any constant $\overline{X} > 0$,*

$$Q(\exists t \ X_t \geq \overline{X}) \leq \frac{X_0}{\overline{X}}.$$

Proof. Let τ be the stopping time when X_t first weakly exceeds $\overline{X}$. If X_t never weakly exceeds $\overline{X}$ then set $\tau > T$. By Doob's optimal sampling, we have $X_0 \geq \mathbf{E}X_\tau = \mathbf{E}X_\tau 1_{\tau \leq T} + \mathbf{E}X_\tau 1_{\tau > T} \geq \overline{X}\mathbf{E}1_{\tau \leq T} = \overline{X}Q(\exists t \ X_t \geq \overline{X})$. □

Proof of Proposition 2. Let $(\hat{\mathbf{P}}, \mathbf{\Omega}, \boldsymbol{\mathcal{F}})$ satisfy $\hat{\mathbf{P}}(\mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})) \geq 1 - \delta(z)$.

By definition, $\boldsymbol{\mu} \in \mathbf{\Omega} - \mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \hat{P}^n, c(z))$ if and only if

$$\hat{P}^n \left(\boldsymbol{\mu}^n \notin \mathbf{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \;\middle|\; \boldsymbol{\mu}^n|_s \right) \geq \frac{\varepsilon}{\beta^{1-s} T \overline{v}} - c(z) \quad \text{for some } s.$$

Define the martingale $\{X_s\}_{s=1}^T$ on $(\hat{\mathbf{P}}, \mathbf{\Omega}, \boldsymbol{\mathcal{F}})$ where

$$X_s(\boldsymbol{\mu}) = \hat{P}^n \left(\boldsymbol{\mu}^n \notin \mathbf{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \;\middle|\; \boldsymbol{\mu}^n|_s \right).$$

The unconditional expectation of this martingale is $\hat{P}^n \left(\boldsymbol{\mu}^n \notin \mathbf{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \right)$. Since $\boldsymbol{\mu}^n \notin \mathbf{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \Rightarrow \boldsymbol{\mu} \notin \mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, it must be that $\hat{\mathbf{P}}(\boldsymbol{\mu} \notin \mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})) \geq \hat{P}^n \left(\boldsymbol{\mu}^n \notin \mathbf{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \right)$. And since $\hat{\mathbf{P}}(\mathbf{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})) \geq 1 - \delta(z)$, we have that $\delta(z) \geq \hat{P}^n \left(\boldsymbol{\mu}^n \notin \mathbf{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \right)$.

Now, by Lemma 5,

$$
\begin{aligned}
\hat{\mathbf{P}}(\boldsymbol{\Omega} - \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \hat{P}^n, c(z))) &\leq \frac{\frac{z(1-z)\varepsilon^2}{(\beta^{1-T}T\overline{v})^2 N} \wedge 1}{\frac{\varepsilon}{\beta^{1-T}T\overline{v}} - c(z)} \\
&\leq \frac{\frac{z(1-z)\varepsilon^2}{(\beta^{1-T}T\overline{v})^2 N}}{\frac{(1-z)\varepsilon}{\beta^{1-T}T\overline{v}}} \\
&\leq \frac{z\varepsilon}{\beta^{1-T}T\overline{v}N}.
\end{aligned}
$$

By De Morgan's Law, $\hat{\mathbf{P}}(\boldsymbol{\Omega} - \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \hat{\mathbf{P}}, c(z)\mathbf{1})) \leq \left(N \cdot \frac{z\varepsilon}{\beta^{1-T}T\overline{v}N}\right) \wedge 1 = c(z)$. □

The justification of reasonable strategies is based on simultaneously assuming agent n faces at least a small cost when thinking about other agents, but faces no cost when updating P^n with respect to $\boldsymbol{\mu}|_t$. This is a rather stark contrast, and one may also wish to introduce some friction to how P^n updates with respect to $\boldsymbol{\mu}|_t$. As long as the frictional updating rule – let us call it $P^n_{fr}(\cdot \mid \boldsymbol{\mu}^n|_s)$ – preserves the martingale property – formally, as long as $\{P^n_{fr}(\boldsymbol{\mu}^n \notin \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \mid \boldsymbol{\mu}^n|_s)\}$ is a martingale while the frictional version of (2) where $P^n(\cdot \mid \boldsymbol{\mu}^n|_s)$ is replaced with $P^n_{fr}(\cdot \mid \boldsymbol{\mu}^n|_s)$ has not been violated – then Proposition 2 and its proof remain valid.

The ideas of this section can be naturally extended to the case $N = 1$. Since there are no other agents to think about, the natural definition of a reasonable agent 1 strategy becomes: $\sigma^1_t(\boldsymbol{\mu}^1|_t) = \mu^1_t$ whenever

$$
\left[\mathbf{P}\left(\boldsymbol{\mu} \notin \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}) \;\middle|\; \boldsymbol{\mu}|_s\right)\right] \beta^{1-s}T\overline{v} < \varepsilon
$$

for all $s \leq t$. The justification for reasonable strategies becomes stronger and resembles that for contemporaneous perfect ε-equilibrium in a one player extensive-form game. See Mailath, Postlewaite, and Samuelson (2005). The analogue of $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \mathbf{P}, \boldsymbol{c})$ is the event – call it $\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \mathbf{P})$ – in which agent 1 employs the truth-telling strategy under any reasonable strategy. An analogue of Proposition 2 is the claim that for any $c \in (0, 1]$, if $\hat{\mathbf{P}}$ is any probability on $(\boldsymbol{\Omega}, \boldsymbol{\mathcal{F}})$ satisfying $\hat{\mathbf{P}}(\boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})) \geq 1 - \delta(c)$, where $\delta(c) := \frac{c\varepsilon}{\beta^{1-T}T\overline{v}} \wedge 1$, then $\hat{\mathbf{P}}(\boldsymbol{\Omega} - \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}}, \hat{\mathbf{P}})) \leq c$. The claim admits the same interpretation as Proposition 2, and, just like in the $N \geq 2$ case, is true because of Lemma 5.

6 The Many Small Agents Case

This section culminates in the construction of a sequence of true models that can be viewed as a single true model scaled to increasingly larger sizes, with the number of agents and the available quantity of each of the R resources increasing proportionally. Theorem 2 then shows that if agents are patient and it is common knowledge something akin to, but weaker than, the Central Limit Theorem holds across agents and across dates, then in the limit, an appropriately linked VCG mechanism achieves allocative efficiency at a cost-to-surplus ratio of zero.

A key step in establishing Theorem 2 will be to show that in the limit model, under the limit linked VCG mechanism, each agent is "small" in a sense to be made precise below:

6.1 The Cost Savings of Linking when an Agent is Small

Fix an arbitrary true model, linked VCG mechanism $(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$, and $\boldsymbol{\mu} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$. Let $\boldsymbol{p}(\boldsymbol{\mu})$ denote the set of T-row vectors of prices associated with the linear investments made under ${}_l\boldsymbol{A}(\boldsymbol{\mu})$, and assume $|\boldsymbol{p}(\boldsymbol{\mu})| = 1$ – that is, the associated price each date is unique.

Consider the following two properties that an agent n might satisfy:

I. $\boldsymbol{p}(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n) = \boldsymbol{p}(\boldsymbol{\mu})$ for all $\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$.

II. $\forall t\ \forall i,\ N(\boldsymbol{\mu}_t^{-n}, \hat{\mu}_t^n, i) > n\ \forall \hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ or $N(\boldsymbol{\mu}_t^{-n}, \hat{\mu}_t^n, i) < n\ \forall \hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$.

An agent who satisfies Property I., especially one who also satisfies Property II., can be described as "taking prices as given." Both properties capture a notion of smallness: By Lemma 2, an agent who has a sufficiently small effect on the efficient allocation will satisfy Property I. As for Property II., note that when N is much larger than TI, a typical agent's index n will be far away from $N(\boldsymbol{\mu}_t, i)$ for all t and i. If this agent has a sufficiently small effect on the efficient allocation, then $N(\boldsymbol{\mu}_t^{-n}, \hat{\boldsymbol{\mu}}^n, i)$ will be close enough to $N(\boldsymbol{\mu}_t, i)$ for all t, i and $\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ that he will satisfy Property II.

In Subsection 6.3, when I consider a limit model with many agents, it will be the case that, under the limit linked VCG mechanism, almost surely, all agents satisfy Property I. and almost all agents also satisfy Property II.

Lemma 6. *If Property I. is satisfied, then*

$$ {}_lW^n(\boldsymbol{\mu}) \le \max_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})} S^n(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n). $$

If, in addition, Property II. is also satisfied, then

$$ {}_lW^n(\boldsymbol{\mu}) = \max_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})} \sum_{t=1}^{T} \beta^{t-T} p_t(\boldsymbol{\mu}) \cdot \left[\sum_{i=1}^{I} \left(\hat{\mu}_{ti}^n - \mu_{ti}^n\right) \theta_i 1_{n > N(\boldsymbol{\mu}_t, i)} \right]. $$

Proof. See appendix. □

Corollary 2. *Suppose $\beta = 1$ and $\mathbf{p}(\boldsymbol{\mu}) = p\mathbf{1}$ for some p. If Properties I. and II. are satisfied, then*

$$ {}_lW^n(\boldsymbol{\mu}) \leq \max_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})} S^n(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n) - \min_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})} S^n(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n). $$

Proof. By Lemma 6,

$$ \begin{aligned} {}_lW^n(\boldsymbol{\mu}) =& p \cdot \left[\sum_{i=1}^{I} \left(T\overline{\mu}^n_{ti} - \sum_{t=1}^{T} \mu^n_{ti} \right) \theta_i 1_{N(\boldsymbol{\mu}_t,i)>n} \right] \\ \leq& \left[\sum_{i=1}^{I} \left(T\overline{\mu}^n_{ti} - \sum_{t=1}^{T} \mu^n_{ti} \right) 1_{N(\boldsymbol{\mu}_t,i)>n} \right] \\ \leq& \sum_{i=1}^{I} \left(T\overline{\mu}^n_{ti} - T\underline{\mu}^n_{ti} \right) 1_{N(\boldsymbol{\mu}_t,i)>n} \\ =& \max_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})} S^n(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n) - \min_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})} S^n(\boldsymbol{\mu}^{-n}, \hat{\boldsymbol{\mu}}^n). \end{aligned} $$

□

6.2 The Geometry of Linking when an Agent is Small

In this subsection I show how the transfer to an agent satisfying Properties I. and II. can be computed via a simple geometric algorithm.

Define the date-valued correspondence $\tau : \Theta \to \{1, \dots T\}$ as follows:

$$ \tau(\theta_i) = \operatorname*{arg\,max}_{t \text{ s.t. } n>N(\boldsymbol{\mu}_t,i)} \beta^{t-T} p_t(\boldsymbol{\mu}) \cdot \theta_i. $$

If the set of times such that $n > N(\boldsymbol{\mu}_t, i)$ is empty, then set $\tau(\theta_i) = \{1, \dots T\}$. Given any selection $\tilde{\tau}$ of τ, define $\tilde{\boldsymbol{\mu}}^n$ as follows: For each t and i, $\tilde{\mu}^n_{ti} = T\overline{\mu}^n_{ti}$ if $\tilde{\tau}(\theta_i) = t$ and $= 0$ otherwise.

It is evident

$$ \tilde{\boldsymbol{\mu}}^n \in \operatorname*{arg\,max}_{\hat{\boldsymbol{\mu}}^n \in \boldsymbol{\Omega}^n(T,\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}})} \sum_{t=1}^{T} \beta^{t-T} p_t(\boldsymbol{\mu}) \cdot \left[\sum_{i=1}^{I} \left(\hat{\mu}^n_{ti} - \mu^n_{ti} \right) \theta_i 1_{n>N(\boldsymbol{\mu}_t,i)} \right]. $$

Thus, given Lemma 6, the problem of computing ${}_lW^n(\boldsymbol{\mu})$ when both Properties I. and II. are satisfied reduces to characterizing $\tau(\cdot)$.

Algorithm

1. Associate to each price $p_t(\boldsymbol{\mu})$, the hyperplane $H_t := \{x \in \mathbb{R}^R \mid p_t(\boldsymbol{\mu}) \cdot x = 1\}$. H_t is said to enclose θ_i if $n > N(\boldsymbol{\mu}_t, i)$. Geometrically, if H_t encloses θ_i, then θ_i

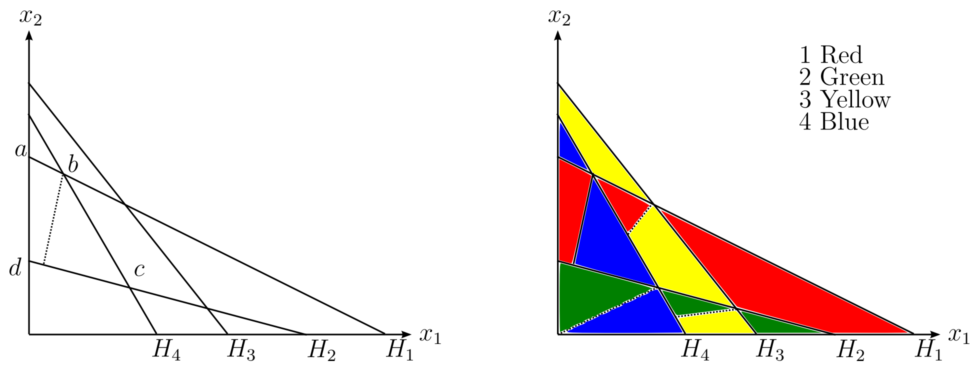

(a) Model with 2 resources and 4 dates.

(b) $\tau(\cdot)$ when $\beta = 1$.

must lie on or below H_t.

Example: Figure 4a depicts a model with four hyperplanes.

2. If $\beta = 1$, then

   For each θ_i, consider the ray going through θ_i. Let $\{H_t\}_{t \in T_i}$ be the set of enclosing hyperplanes that the ray touches first. If $T_i = \emptyset$ then set $\tau(\theta_i) = \{1, \ldots T\}$. Otherwise, set $\tau(\theta_i) = T_i$.

   Example: In Figure 4a, the dotted line is part of the ray that goes through $H_1 \cap H_4$. For any $\theta_i \in Int(abcd)$, the enclosing hyperplanes are $\{H_1, H_3, H_4\}$. If θ_i is to the left of the dotted line, then $\tau(\theta_i) = 1$. If θ_i is on the dotted line, then $\tau(\theta_i) = \{1, 4\}$. If θ_i is to the right of the dotted line, then $\tau(\theta_i) = 4$.

   Figure 4b color codes everywhere $\tau(\theta_i)$ takes on a single value.

3. If $\beta < 1$, then

   For each θ_i, and each of its enclosing H_t, create $\tilde{H}_t := \{x \in \mathbb{R}^R \mid \beta^{t-1} p_t(\boldsymbol{\mu}) \cdot x = 1\}$. Define $\tau(\theta_i)$ by redoing Step 2. with $\{\tilde{H}_t\}$ replacing the set of enclosing hyperplanes.

6.3 A Cheap Efficient Linked VCG Mechanism

Consider the following sequence of true models, parameterized by $k = 1, 2, 3, \ldots$, that can be viewed as a single true model scaled to increasingly larger sizes.

For each k, I will append $[k]$ to parameters in the k-model. Suppose there exist Θ, $\overline{q}$, δ and Δ such that, for all k, $\beta[k] = 1$, $N[k] = k$, $\Theta[k] = \Theta$, and $\overline{q}[k] = k\overline{q}$. Also, suppose for any $k_1 \leq k_2$, $n \leq k_1$, and t, we have $\mu_t^n[k_1] = \mu_t^n[k_2]$ – this means the vector of k_1-model endowment processes can be identified as a subset of the vector of k_2-model endowment processes. Consequently, I do not append $[k]$ to μ_t^n from now on and, as an abuse of notation, I let $\mathbf{P}$ denote the governing probability of the universal set of endowments $\{\mu_t^n\}_{1 \leq n, t, < \infty}$ as well as the governing probability of any k-model.

Once again, $\mathbf{P}$ need not be common knowledge. However,

Assumption 1. *There exist*

- $\underline{\pi}, \overline{\pi}, \{\pi^n\}_{n=1}^{\infty} \in \Omega$, *satisfying* $|p(\overline{q}, \pi + \mu_t^0)| = 1$, $\pi^n \leq \overline{\pi}\ \forall n$,
- $\kappa \in (0,1)$ *and weakly decreasing functions* $f, F : [0,\infty) \to [0,\infty)$ *satisfying* $\lim_{z\to\infty} z^{\frac{3}{\kappa}} f(z) = z^{\kappa+1}F(z) = 0$, *and*
- *weakly increasing functions* $i, I : [0,\infty) \to [0,\infty)$,

such that, for all k*, it is common knowledge in the* k*-model that*

$$P^n\left[\left\|\frac{\sum_{t=1}^{T}\mu_t^n}{T} - \pi^n\right\|_\infty \leq d\right] \geq 1 - f(i(d)T) \qquad \forall d > 0, n \leq k, T, \tag{3}$$

$$\mathbf{P}\left[\left\|\frac{\sum_{n=1}^{k}\mu_t^n}{k} - \pi\right\|_\infty \leq d\right] \geq 1 - F(I(d)k) \qquad \forall d > 0, t. \tag{4}$$

Assumption 1 requires that the strength of the ergodicity of endowments across dates be uniformly bounded below across agents and that the strength of the ergodicity of endowments across agents be uniformly bounded below across dates. The condition $\lim_{z\to\infty} z^{\frac{3}{\kappa}} f(z) = z^{\kappa+1}F(z) = 0$ means that the uniform bounds cannot be too weak. For example, if $\mathbf{P}$ is common knowledge and, under $\mathbf{P}$, endowments are iid across dates and agents with finite first and second moments, then the Central Limit Theorem implies that $f(z)$ and $F(z)$ can both be defined to be $\exp(-z)$, in which case, the condition is satisfied for any choice of κ.

Theorem 2. *There exists a sequence of linked VCG mechanisms* $(T[k], \underline{\boldsymbol{\mu}}[k], \overline{\boldsymbol{\mu}}[k])$ *and conjectures* $c[k]$*, one for each* k*-model, such that* $c[k]\mathbf{1}$ *is believable in the* k*-model for all* k*,* $\lim_{k\to\infty} c[k] = 0$*, and*

$$\lim_{k\to\infty} \sup_{\text{reasonable } \boldsymbol{\sigma} \text{ given } c[k]\mathbf{1}} \frac{\mathbf{E_P}\ {}_{\iota}C(\boldsymbol{\sigma}(\boldsymbol{\mu}))}{\mathbf{E_P}\ {}_{\iota}S(\boldsymbol{\sigma}(\boldsymbol{\mu}))} = 0.$$

Theorem 2 says that, in the limit, the principal can, under an appropriately linked VCG mechanism, expect to achieve allocative efficiency at a cost-to-surplus ratio of zero.

The proof of the zero cost-to-surplus ratio result is organized as follows: As $k \to \infty$, the probability of the event that all agents employ the truth-telling strategy under any reasonable strategy profile given believable $c[k]\mathbf{1}$ converges to 1. In the limit, as this increasingly sure event becomes almost sure, Assumption 1 implies $\boldsymbol{p}(\boldsymbol{\mu}) = p(\overline{q}, \pi)\mathbf{1}$, all agents satisfy Property I., and almost all agents also satisfy Property II. Corollary 2 implies that those agents who satisfy both properties are paid at most the difference between their maximum and minimum contributions to surplus. Since the sequence of linked VCG mechanisms I will create satisfies $\lim_{k\to\infty} \overline{\boldsymbol{\mu}}[k] - \underline{\boldsymbol{\mu}}[k] = 0$, it will be the case that the amount paid to agents who satisfy both properties comprise a zero fraction of surplus in the limit. The remaining fraction of agents – which will be zero in the limit – still satisfy Property I. Therefore, by Lemma 6, they are paid at most their maximum contributions to surplus. Again, since $\lim_{k\to\infty} \overline{\boldsymbol{\mu}}[k] - \underline{\boldsymbol{\mu}}[k] = 0$, these agents' maximum contributions to surplus are close to their actual contributions to surplus. Thus, the amount paid to agents who just satisfy Property I. also comprise a zero fraction of surplus in the limit. These two observations – plus some fine tuning of the sequence of linked VCG mechanisms to make sure that the aforementioned increasingly sure event converges in probability to 1 sufficiently quickly – imply that the cost-to-surplus ratio of the limiting linked VCG mechanism is zero.

Proof. Fix a κ satisfying Assumption 1 and let $T[k] = \lceil k^\kappa \rceil$. For each $d > 0$, define

$$c(d)[k] := \frac{\frac{\varepsilon}{k\overline{v}} - \sqrt{\left(\frac{\varepsilon}{k\overline{v}}\right)^2 - 4kf(i(d)T[k])}}{2}$$

whenever the right hand side $\in [0,1]$, which is true for all sufficiently large k.[2] Otherwise, set $c(d)[k] = 1$.

For each $d > 0$ and n, define $\underline{\mu}^n(d)$ and $\overline{\mu}^n(d)$ to be the elements of Ω satisfying $\underline{\mu}^n(d)_i = \max\{\pi_i^n - d, 0\}$ and $\overline{\mu}^n(d)_i = \pi_i^n + d$ for all i. Define $\underline{\boldsymbol{\mu}}(d)[k]$ and $\overline{\boldsymbol{\mu}}(d)[k]$ accordingly.

Let $\hat{\mathbf{P}}$ be any probability on $(\boldsymbol{\Omega}, \boldsymbol{\mathcal{F}})$ satisfying Assumption 1. Lemma 5 and the quadratic formula imply that any

$$\hat{\mathbf{P}}(\boldsymbol{\Omega} - \boldsymbol{\Omega}(T[k], \underline{\boldsymbol{\mu}}(d)[k], \overline{\boldsymbol{\mu}}(d)[k], \hat{\mathbf{P}}, c(d)[k]\mathbf{1})) \leq c(d)[k]$$

in the k-model.[3]

[2] Let us prove the stronger result: For all sufficiently large k, $k \cdot c(d)[k] \in [0, \infty)$ and $\lim_{k\to\infty} k \cdot c(d)[k] = 0$. We have

$$k \cdot c(d)[k] = \frac{\frac{\varepsilon}{\overline{v}} - \sqrt{\left(\frac{\varepsilon}{\overline{v}}\right)^2 - 4k^3 f(i(d)T[k])}}{2}.$$

It suffices to show $\lim_{k\to\infty} k^3 f(i(d)T[k]) = 0$, which follows from Assumption 1.

[3] It suffices to prove this claim when $c(d)[k] < 1$. Note, by De Morgan's Law,

$$\hat{\mathbf{P}}(\boldsymbol{\Omega} - \boldsymbol{\Omega}(T[k], \underline{\boldsymbol{\mu}}(d)[k], \overline{\boldsymbol{\mu}}(d)[k], \hat{\mathbf{P}}, c(d)[k]\mathbf{1})) \leq \sum_{n=1}^{k} \hat{\mathbf{P}}(\boldsymbol{\Omega} - \boldsymbol{\Omega}^n(T[k], \underline{\mu}^n(d), \overline{\mu}^n(d), \hat{P}^n, c(d)[k])).$$

Mirroring the proof of Proposition 2, note that $\boldsymbol{\mu} \in \boldsymbol{\Omega} - \boldsymbol{\Omega}^n(T[k], \underline{\mu}^n(d), \overline{\mu}^n(d), \hat{P}^n, c(d)[k])$ if and only if

$$\hat{P}^n\left(\boldsymbol{\mu}^n \notin \boldsymbol{\Omega}^n(T[k], \underline{\mu}^n(d), \overline{\mu}^n(d)) \mid \boldsymbol{\mu}^n|_s\right) \geq \frac{\varepsilon}{T[k]\overline{v}} - c(d)[k] \qquad \text{for some } s \leq T[k].$$

Define $\mathbf{\Omega}^*(d)[k]$ to be the set of $\boldsymbol{\mu} \in \mathbf{\Omega}(T[k], \underline{\boldsymbol{\mu}}(d)[k], \overline{\boldsymbol{\mu}}(d)[k], \mathbf{P}, c(d)[k]\mathbf{1})$ satisfying

$$\left\| \frac{\sum_{n=1}^{k} \mu_t^n}{k} - \pi \right\|_\infty \leq d \quad \forall t \leq T[k].$$

Assumption 1 and De Morgan's Law imply

$$\mathbf{P}(\mathbf{\Omega} - \mathbf{\Omega}^*(d)[k]) \leq c(d)[k] + k^\kappa F(I(d)k).$$

Since Assumption 1 implies $\lim_{k\to\infty} k \cdot k^\kappa F(I(d)k) = 0$ and a previous footnote showed that $\lim_{k\to\infty} k \cdot c(d)[k] = 0$, therefore

$$\lim_{k\to\infty} k\left(c(d)[k] + k^\kappa F(I(d)k)\right) = 0 \quad \forall d > 0.$$

This means it is possible to pick a decreasing sequence $\{d_k\}_{k=1,2,\ldots}$ converging to 0 such that $\lim_{k\to\infty} k(c(d_k)[k] + k^\kappa F(I(d_k)k)) = 0$.

Define $c[k] = c(d_k)[k]$, $\underline{\boldsymbol{\mu}}[k] = \underline{\boldsymbol{\mu}}(d_k)[k]$, $\overline{\boldsymbol{\mu}}[k] = \overline{\boldsymbol{\mu}}(d_k)[k]$ and $\mathbf{\Omega}^*[k] = \mathbf{\Omega}^*(d_k)[k]$. Based on what we have already shown, the first two parts of the theorem have been proved. It remains to prove the last part, concerning the cost-to-surplus ratio.

For any $k, n \leq k, t \leq T[k]$, and $\hat{\boldsymbol{\mu}}^n \in \mathbf{\Omega}^n(T[k], \underline{\boldsymbol{\mu}}[k], \overline{\boldsymbol{\mu}}[k])$, we have $\|\hat{\mu}_t^n\|_\infty \leq \|\hat{\boldsymbol{\mu}}^n \mathbf{1}\|_\infty \leq k^\kappa(\|\pi^n\|_\infty + d_k) \leq k^\kappa(\|\overline{\pi}\|_\infty + d_k)$.

By Corollary 1, for any k, $\boldsymbol{\mu} \in \mathbf{\Omega}^*[k]$, agent $n \leq k$ and date t, we have

$$p\left(k\overline{q}, \sum_{m\neq n, m=1}^{k} \mu_t^n + \hat{\mu}_t^n + \mu_t^0\right) = p\left(\overline{q}, \frac{\sum_{m\neq n,m=1}^{k} \mu_t^n + \hat{\mu}_t^n}{k} + \mu_t^0\right).$$

Moreover,

$$\begin{aligned} \left\| \frac{\sum_{m\neq n,m=1}^{k} \mu_t^n + \hat{\mu}_t^n}{k} + \mu_t^0 - (\pi + \mu_t^0) \right\|_\infty &\leq \left\| \frac{\sum_{m=1}^{k} \mu_t^n}{k} - \pi \right\|_\infty + \left\| \frac{\mu_t^n}{k} \right\|_\infty + \left\| \frac{\hat{\mu}_t^n}{k} \right\|_\infty \\ &\leq d_k + 2\frac{\|\overline{\pi}\|_\infty + d_k}{k^{1-\kappa}}. \end{aligned}$$

This last quantity goes to zero as $k \to \infty$. Lemma 2 now implies for all k sufficiently large and all $\boldsymbol{\mu} \in \mathbf{\Omega}^*[k]$, all agents $n \leq k$ satisfy Property I with the unique price

---

Also, by Assumption 1, the unconditional probability, $\hat{P}^n\left(\boldsymbol{\mu}^n \notin \mathbf{\Omega}^n(T[k], \underline{\mu}^n(d), \overline{\mu}^n(d)\right)$, is less than or equal to $f(i(d)T[k])$. Thus, by Lemma 5, we have

$$\hat{P}^n(\mathbf{\Omega} - \mathbf{\Omega}^n(T[k], \underline{\mu}^n(d), \overline{\mu}^n(d), \hat{P}^n, c(d)[k])) \leq \frac{f(i(d)T[k])}{\left(\frac{\varepsilon}{T[k]\overline{v}} - c(d)[k]\right)} \leq \frac{f(i(d)T[k])}{\left(\frac{\varepsilon}{k\overline{v}} - c(d)[k]\right)}.$$

Thus, $\hat{\mathbf{P}}(\mathbf{\Omega} - \mathbf{\Omega}(T[k], \underline{\boldsymbol{\mu}}(d)[k], \overline{\boldsymbol{\mu}}(d)[k], \hat{\mathbf{P}}, c(d)[k]\mathbf{1})) \leq \frac{kf(i(d)T[k])}{\left(\frac{\varepsilon}{k\overline{v}} - c(d)[k]\right)} = c(d)[k]$.

each date being $p := p(\overline{q}, \pi + \mu_t^0)$.

For each k, let $N^*[k]$ be a lower bound on the number of agents $n \leq k$ for whom Property II. is satisfied across all $\boldsymbol{\mu} \in \boldsymbol{\Omega}^*[k]$. A similar argument implies $N^*[k]$ can be chosen so that $\lim_{k\to\infty} N^*[k]/k = 1$.

Now, by Corollary 2 and Lemma 6, for all k sufficiently large, when $\boldsymbol{\mu} \in \boldsymbol{\Omega}^*[k]$, the transfer to agent $n \leq k$ is bounded above by $k^\kappa 2 d_k I$ if Properties I. and II. are satisfied, and bounded above by $k^\kappa \sum_i \overline{\mu}_i^n[k] \leq k^\kappa(\|\overline{\pi}\|_\infty + d_k)I$ if only Property I. is satisfied. Thus, the total transfer, when $\boldsymbol{\mu} \in \boldsymbol{\Omega}^*[k]$, is bounded above by

$$[N^*[k]k^\kappa 2d_k + (k - N^*[k])k^\kappa(\|\overline{\pi}\|_\infty + d_k)] \cdot I.$$

In general, the transfer to agent n is bounded above by $k^\kappa k \overline{v}$. Putting everything together, we have that

$$\sup_{\text{reasonable } \boldsymbol{\sigma} \text{ given } c[k]\mathbf{1}} \mathbf{E}_{\mathbf{P}}\ {}_{\iota}C(\boldsymbol{\sigma}(\boldsymbol{\mu})) \leq (c(d_k)[k] + k^\kappa F(I(d_k)k))\, k k^\kappa k \overline{v} + [N^*[k]k^\kappa 2d_k + (k - N^*[k])k^\kappa(\|\overline{\pi}\|_\infty + d_k)] \cdot I.$$

On the other hand,

$$\inf_{\text{reasonable } \boldsymbol{\sigma} \text{ given } c[k]\mathbf{1}} \mathbf{E}_{\mathbf{P}}\ {}_{\iota}S(\boldsymbol{\sigma}(\boldsymbol{\mu})) \geq \mathbf{P}(\boldsymbol{\Omega}^*[k]) k^\kappa k \underline{v},$$

where $\underline{v} = v(\overline{q}, \mu_t^0)$. Thus,

$$\begin{aligned}
&\lim_{k\to\infty} \sup_{\text{reasonable } \boldsymbol{\sigma} \text{ given } c[k]\mathbf{1}} \frac{\mathbf{E}_{\mathbf{P}}\ {}_{\iota}C(\boldsymbol{\sigma}(\boldsymbol{\mu}))}{\mathbf{E}_{\mathbf{P}}\ {}_{\iota}S(\boldsymbol{\sigma}(\boldsymbol{\mu}))} \\
&\lim_{k\to\infty} \frac{\sup_{\text{reasonable } \boldsymbol{\sigma} \text{ given } c[k]\mathbf{1}} \mathbf{E}_{\mathbf{P}}\ {}_{\iota}C(\boldsymbol{\sigma}(\boldsymbol{\mu}))}{\inf_{\text{reasonable } \boldsymbol{\sigma} \text{ given } c[k]\mathbf{1}} \mathbf{E}_{\mathbf{P}}\ {}_{\iota}S(\boldsymbol{\sigma}(\boldsymbol{\mu}))} \\
&\leq \lim_{k\to\infty} \left\{ (c(d_k)[k] + k^\kappa F(I(d_k)k))\, k \frac{\overline{v}}{\mathbf{P}(\boldsymbol{\Omega}^*[k])\underline{v}} + \right. \\
&\quad \left. \frac{N^*[k]}{k} \frac{2d_k I}{\mathbf{P}(\boldsymbol{\Omega}^*[k])\underline{v}} + \left(1 - \frac{N^*[k]}{k}\right) \frac{(\|\overline{\pi}\|_\infty + d_k)I}{\mathbf{P}(\boldsymbol{\Omega}^*[k])\underline{v}} \right\} \\
&= 0.
\end{aligned}$$

□

7 Appendix

Proof of Lemma 1. Uniqueness of $\alpha(q, \mu)$. Let $\alpha^*, \tilde{\alpha}$ be two exhaustive and linear investments.

Case 1. α^* and $\tilde{\alpha}$ have a common associated price p^*. Let H_{p^*} denote the hyperplane $\{x \in \mathbb{R}^R \mid p^* \cdot x = 1\}$.

Since α^* and $\tilde{\alpha}$ are both exhaustive and linear,

$$\sum_{\theta_i \in H_{p^*}} \alpha_i^* \theta_i = q - \sum_{p^* \cdot \theta_i < 1} \mu_i \theta_i = \sum_{\theta_i \in H_{p^*}} \tilde{\alpha}_i \theta_i.$$

Since $\{\theta_i \in H_{p^*}\}$ is linearly independent, $\sum_{\theta_i \in H_{p^*}} \alpha_i^* \theta_i = \sum_{\theta_i \in H_{p^*}} \tilde{\alpha}_i \theta_i \Rightarrow \alpha_i^* = \tilde{\alpha}_i$ for $\theta_i \in H_{p^*} \Rightarrow \alpha^* = \tilde{\alpha}$.

Case 2. α^* and $\tilde{\alpha}$ have associated prices $p^* \neq \tilde{p}$. If $p^*(r) < \tilde{p}(r)$ for all r, then linearity implies $\alpha_i^* \geq \tilde{\alpha}_i$ for all i, and then exhaustiveness implies $\alpha^* = \tilde{\alpha}$. Similarly, if $p^*(r) > \tilde{p}(r)$ for all r then $\alpha^* = \tilde{\alpha}$.

Otherwise, there exist r and r', not necessarily distinct, such that $p^*(r) \geq \tilde{p}(r)$ and $p^*(r') \leq \tilde{p}(r')$. This implies the set $X = \{x \in [0,\infty)^R - \{0\} \mid p^* \cdot x = \tilde{p} \cdot x = 1\}$ is non-empty. In addition, define two more possibly empty convex sets

$$X^* = \{x \in \mathbb{R}_+^R \mid p^* \cdot x \leq 1 \leq \tilde{p} \cdot x,\ p^* \cdot x \neq \tilde{p} \cdot x\},$$
$$\tilde{X} = \{x \in \mathbb{R}_+^R \mid p^* \cdot x \geq 1 \geq \tilde{p} \cdot x,\ p^* \cdot x \neq \tilde{p} \cdot x\}.$$

By linearity,

$$\alpha_i^* - \tilde{\alpha}_i = 0 \text{ if } \theta_i \notin X \cup X^* \cup \tilde{X},$$
$$\alpha_i^* - \tilde{\alpha}_i \geq 0 \text{ if } \theta_i \in X^*,$$
$$\alpha_i^* - \tilde{\alpha}_i \leq 0 \text{ if } \theta_i \in \tilde{X}.$$

If $\alpha_i^* - \tilde{\alpha}_i = 0$ for all $\theta_i \in X^* \cup \tilde{X}$, then there exists a common associated price, and $\alpha^* = \tilde{\alpha}$. So suppose there exists at least one $\theta_i \in X^*$ such that $\alpha_i^* - \tilde{\alpha}_i > 0$ or at least one $\theta_i \in \tilde{X}$ such that $\alpha_i^* - \tilde{\alpha}_i < 0$.

There is a unique hyperplane H going through the origin and $H_{p^*} \cap H_{\tilde{p}}$. $X \subset H$, and X^* and $\tilde{X}$ lie on opposite sides of H. Let z be a normal vector of H on the side of X^*. Then $z \cdot \theta_i = 0$ for all $\theta_i \in X$, and $z \cdot \theta_i > 0$ for all $\theta_i \in X^*$ and $z \cdot \theta_i < 0$ for all $\theta_i \in \tilde{X}$. This implies $z \cdot \sum_i (\alpha_i^* - \tilde{\alpha}_i)\theta_i > 0$ which implies $\sum_i (\alpha_i^* - \tilde{\alpha}_i)\theta_i \neq 0$. But $0 = q - q = \sum_i (\alpha_i^* - \tilde{\alpha}_i)\theta_i$. Contradiction.

Existence of $\alpha(q,\mu)$. It suffices to show that an efficient investment exists and that it must be exhaustive and linear.

Since (1) is the maximization of a continuous function over a compact set, an efficient investment α^* exists. Since $\mu \geq \mu_t^0$, α^* must be exhaustive.

Given a set $S \in \mathbb{R}_+^R$, define $Conv(S)$ to be the convex hull of S, $Int(S)$ to be the interior of S, and $Gtr(S) := \{x \in \mathbb{R}^R \mid \exists s \in S,\ x \geq s\}$.

Claim: $Int(Gtr(Conv(Supp(\mu - \alpha^*)))) \cap Conv(Supp(\alpha^*)) = \emptyset$.

Proof of Claim. Suppose not. Let $\tilde{x}$ be an element of the intersection.

Since $\tilde{x} \in Conv(Supp(\alpha^*))$, there exist $x_1, \ldots x_d \in Supp(\alpha^*)$ and $\lambda_1, \ldots \lambda_d \in (0,1)$ summing to one, such that $\tilde{x} = \lambda_1 x_1 + \ldots + \lambda_d x_d$.

By assumption, there exists $\tilde{y} \in Gtr(Conv(Supp(\mu - \alpha^*)))$ such that $\tilde{y}(r) < \tilde{x}(r)$ for all r. By definition of Gtr, one can choose $\tilde{y}$ to be in $Conv(Supp(\mu - \alpha^*))$. Thus, there exist $y_1, \ldots y_e \in Supp(\mu - \alpha^*)$ and $l_1, \ldots l_e \in (0,1)$ summing to one, such that $\tilde{y} = l_1 y_1 + \ldots + l_e y_e$.

Given $\varepsilon > 0$, add an εl_i amount of mass from α^* at y_j, for $j = 1, \ldots e$, and, at the same time, take an $\varepsilon \lambda_j$ amount of mass to α^* at x_j, for $j = 1, \ldots d$. Call it α'. If ε is sufficiently small, $\alpha' \in [0, \mu]$. By construction, $\sum_i \alpha'_i = \sum_i \alpha^*_i$ and $\sum_i \alpha'_i \theta_i(r) < \sum_i \alpha^*_i \theta_i(r) \leq q(r)$ for all r. Now pick an arbitrary y_j and add back a $\delta > 0$ amount of mass to α' at y_j. Call it α''. If δ is sufficiently small, $\sum_i \alpha''_i \theta_i \leq q$. But $\sum_i \alpha''_i > \sum_i \alpha^*_i$, contradicting the efficiency of α^*. □

Now, by the Separating Hyperplane Theorem, there exists a $p(\alpha^*) \in \mathbb{R}^R$ such that $p(\alpha^*) \cdot Conv(Supp(\alpha^*)) \leq 1 \leq p(\alpha^*) \cdot Int(Gtr(Conv(Supp(\mu - \alpha^*))))$. Since $Gtr(Conv(Supp(\mu - \alpha^*)))$ equals the closure of $Int(Gtr(Conv(Supp(\mu - \alpha^*))))$, we have $p(\alpha^*) \cdot Gtr(Conv(Supp(\mu - \alpha^*))) \geq 1$. This implies $p(\alpha^*) \in (0, \infty)^R$ and $p(\alpha^*) \cdot Supp(\alpha^*) \leq 1 \leq p(\alpha^*) \cdot Supp(\mu - \alpha^*)$. □

Proof of Lemma 2. I start be showing, given $(q, \mu) \in QM$, $|p(q, \mu)| = 1$ if and only if $|\{i \mid \alpha_i(q, \mu) \in (0, \mu_i)\}| = R$. Let Θ^o be the subset of Θ satisfying $\alpha_i(q, \mu) \in (0, \mu_i)$. Since $\Theta^o \subset H_p$ for any $p \in p(q, \mu)$, it is linearly independent, and therefore $|\Theta^o| \leq R$. Suppose $|\Theta^o| = R$. If there are two distinct prices $p, p' \in p(q, \mu)$, then $\Theta^o \subset H_p \cap H_{p'}$. But then Θ^o is not linearly independent. Contradiction. Suppose $|\Theta^o| < R$ and p is an associated price. Then any sufficiently close neighbor of p with hyperplane containing Θ^o will also be an associated price, in which case the associated price is not unique.

Next, I prove $\{(q, \mu) \in QM \mid |p(q, \mu)| = 1\}$ is an open dense subset of QM. Openness will follow from the last part. To prove density, fix a $(q, \mu) \in QM$.

I claim there exists an associated price p such that $|\{\theta_i \in H_p\}| = R$: If p' is an associated price with $|\{\theta_i \in H_{p'}\}| < R$, then all sufficiently close neighbors, p'', of p' satisfying $\{\theta_i \in H_{p'}\} \subset H_{p''}$ are also associated prices. At least one of them will correspond to a hyperplane that contains at least one additional point of Θ. The existence of p follows by induction.

For each $\theta_i \in H_p$, perturb $\alpha(q, \mu)_i$ slightly to some $\tilde{\alpha}_i \in (0, \mu_i)$. Define $\tilde{q} := \sum_{p \cdot \theta_i < 1} \mu_i \theta_i + \sum_{\theta_i \in H_p} \tilde{\alpha}_i \theta_i$. Then p is the unique associated price of $(\tilde{q}, \mu)$. By making the perturbation smaller and smaller, $\tilde{q} \to q$ in the limit.

Suppose $|p(q,\mu)| = 1$. Then the elements of Θ^o span $\mathbb{R}^R$. Thus, given any sufficiently small perturbation of q to $\tilde{q}$ and any sufficiently small perturbation of μ to $\tilde{\mu}$, then there exists a small enough perturbation of $\{\alpha(q,\mu)_i\}_{\theta_i \in \Theta^o}$ to $\{\tilde{\alpha}_i\}_{\theta_i \in \Theta^o}$ such that $\tilde{\alpha}_i \in (0, \tilde{\mu}_i)$ for all $\theta_i \in \Theta^o$, and its extension to the investment $\tilde{\alpha}$ where $\tilde{\alpha}_i = \tilde{\mu}_i (= 0)$ if $p(q,\mu) \cdot \theta_i < 1 (> 1)$ satisfies $\sum_i \tilde{\alpha}_i \theta_i = \tilde{q}$. □

Proof of Lemma 3. Concavity of $v(q,\cdot)$ is obvious. Fix $\mu, \nu \geq \mu_t^0$. For any $\lambda \in (0,1]$, define $\alpha(\lambda)$ and $q(\lambda)$ as follows:

- $p(q,\mu) \cdot \theta_i < 1 \Rightarrow \alpha(\lambda)_i = (1-\lambda)\mu_i + \lambda \mu_i \vee \nu_i$,
  - This case can only occur for θ_i with $i > R$,
- $p(q,\mu) \cdot \theta_i = 1 \Rightarrow \alpha(\lambda)_i = (1-\lambda)\alpha(q,\mu)_i + \lambda \mu_i \vee \nu_i$ if $i > R$ and $\alpha(\lambda)_i = (1-\lambda)\alpha(q,\mu)_i + \lambda \frac{\overline{q}(i)}{\theta_i(i)}$ if $i \leq R$,
- $p(q,\mu) \cdot \theta_i > 1 \Rightarrow \alpha(\lambda)_i = 0$,
- $q(\lambda) = \sum_i \alpha(\lambda)_i \theta_i$.

Lemma 1 implies $\alpha(\lambda)$ is the efficient investment given $(q(\lambda), (1-\lambda)\mu + \lambda \mu \vee \nu)$. It is also clear that $(q(\lambda), (1-\lambda)\mu + \lambda\mu \vee \nu) \geq (q, (1-\lambda)\mu + \lambda\nu)$. Thus,

$$\frac{v(q,(1-\lambda)\mu + \lambda\nu) - v(q,\mu)}{\lambda}$$

$$\leq \frac{v(q(\lambda),(1-\lambda)\mu + \lambda\mu \vee \nu) - v(q,\mu)}{\lambda}$$

$$= \frac{\sum_{p(q,\mu)\cdot\theta_i \leq 1, i>R} \alpha(\lambda)_i - \alpha(q,\mu)_i}{\lambda} + \frac{\sum_{p(q,\mu)\cdot\theta_i = 1, i\leq R} \alpha(\lambda)_i - \alpha(q,\mu)_i}{\lambda}$$

$$= \frac{\sum_{p(q,\mu)\cdot\theta_i \leq 1, i>R} \lambda(\mu_i \vee \nu_i - \alpha(q,\mu_i))}{\lambda} + \frac{\sum_{p(q,\mu)\cdot\theta_i = 1, i\leq R} \lambda\left(\frac{\overline{q}(i)}{\theta_i(i)} - \alpha(q,\mu_i)\right)}{\lambda}$$

$$< \frac{\sum_{p(q,\mu)\cdot\theta_i \leq 1, i>R} \lambda\mu_i \vee \nu_i}{\lambda} + \frac{\sum_{p(q,\mu)\cdot\theta_i = 1, i\leq R} \lambda\frac{\overline{q}(i)}{\theta_i(i)}}{\lambda}$$

$$\leq \sum_{i>R} \mu_i \vee \nu_i + \sum_{i\leq R} \frac{\overline{q}(i)}{\theta_i(i)}.$$

Since this inequality is true for all λ, we must have $\frac{d}{d\lambda}^{+}\big|_{\lambda=0} v(q,(1-\lambda)\mu+\lambda\nu) \leq \sum_{i>R}\mu_i \vee \nu_i + \sum_{i\leq R}\frac{\overline{q}(i)}{\theta_i(i)}$. Concavity of $v(q,\cdot)$ now implies $\frac{d}{d\lambda}^{-}\big|_{\lambda=\lambda^*} v(q,(1-\lambda)\mu+\lambda\nu) \leq \sum_{i>R}\mu_i \vee \nu_i + \sum_{i\leq R}\frac{\overline{q}(i)}{\theta_i(i)}$ for all $\lambda^* \in (0,1]$. □

Proof of Lemma 6. Let $\hat{\boldsymbol{\mu}} \in \boldsymbol{\Omega}(T, \underline{\boldsymbol{\mu}}, \overline{\boldsymbol{\mu}})$ satisfy $\hat{\boldsymbol{\mu}}^{-n} = \boldsymbol{\mu}^{-n}$. By definition,

$$_uW^n(\boldsymbol{\mu}^n) - {_uW^n}(\hat{\boldsymbol{\mu}})$$

$$= \sum_{t=1}^{T} \beta^{t-T} \sum_{m\neq n, m=0}^{N} \sum_{i=1}^{I} \Bigg\{ \mu_{ti}^m 1_{m>N(\mathbf{1}\boldsymbol{\mu}_t,i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) - \sum_{l=m+1}^{N} \mu_{ti}^l \right) 1_{m=N(\mathbf{1}\boldsymbol{\mu}_t,i)}$$

$$- \left[\hat{\mu}_{ti}^m 1_{m>N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\hat{\boldsymbol{\mu}}_t + \mu_t^0) - \sum_{l=m+1}^{N} \hat{\mu}_{ti}^l \right) 1_{m=N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)} \right] \Bigg\}.$$

In the above equation, $_uW^n(\boldsymbol{\mu})$ and $_uW^n(\hat{\boldsymbol{\mu}})$ have both been expressed as sums of the discounted invested project payoffs of everyone except agent n.

Consider a date t project of type i where $p_t(\boldsymbol{\mu}) \cdot \theta_i < 1$. Property I. implies any such project belonging to any agent $m \neq n$ will receive investment and its discounted payoff will be a summand of agent n's unlinked VCG transfer under both $\boldsymbol{\mu}$ and $\hat{\boldsymbol{\mu}}$.

Symmetrically, consider a date t project of type i where $p_t(\boldsymbol{\mu}) \cdot \theta_i > 1$. Property I. implies any such project belonging to any agent $m \neq n$ will not receive investment and its unrealized discounted payoff will not be a summand of agent n's transfer under both $\boldsymbol{\mu}$ and $\hat{\boldsymbol{\mu}}$.

Thus, the above expression can be further simplified by replacing $\sum_{i=1}^{I}$ with $\sum_{\{i \ | \ p_t(\boldsymbol{\mu})\cdot\theta_i = 1\}}$:

$$= \sum_{t=1}^{T} \beta^{t-T} \sum_{m\neq n, m=0}^{N} \sum_{\{i \ | \ p_t(\boldsymbol{\mu})\cdot\theta_i=1\}} \Bigg\{ \mu_{ti}^m 1_{m>N(\mathbf{1}\boldsymbol{\mu}_t,i)} +$$

$$\left(\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) - \sum_{l=m+1}^{N} \mu_{ti}^l \right) 1_{m=N(\mathbf{1}\boldsymbol{\mu}_t,i)}$$

$$- \left[\hat{\mu}_{ti}^m 1_{m>N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\hat{\boldsymbol{\mu}}_t + \mu_t^0) - \sum_{l=m+1}^{N} \hat{\mu}_{ti}^l \right) 1_{m=N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)} \right] \Bigg\}$$

$$= \sum_{t=1}^{T} \beta^{t-T} p_t(\boldsymbol{\mu}) \cdot \sum_{m \neq n, m=0}^{N} \sum_{\{i \mid p_t(\boldsymbol{\mu}) \cdot \theta_i = 1\}} \Bigg\{ \mu_{ti}^m \theta_i 1_{m > N(\mathbf{1}\boldsymbol{\mu}_t, i)} +$$

$$\left(\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) - \sum_{l=m+1}^{N} \mu_{ti}^l \right) \theta_i 1_{m = N(\mathbf{1}\boldsymbol{\mu}_t, i)}$$

$$- \left[\hat{\mu}_{ti}^m \theta_i 1_{m > N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\hat{\boldsymbol{\mu}}_t + \mu_t^0) - \sum_{l=m+1}^{N} \hat{\mu}_{ti}^l \right) \theta_i 1_{m = N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} \right] \Bigg\}.$$

Notice, in the manipulation above, I have rewritten what is inside $\sum_{m \neq n, m=0}^{N}$ from a summation of project payoffs to a summation of resources and factored out the price vector $p_t(\boldsymbol{\mu})$.

By Lemma 1, we know efficient investment is exhaustive. This means, for each t,

$$\sum_{m \neq n, m=0}^{N} \sum_{\{i \mid p_t(\boldsymbol{\mu}) \cdot \theta_i = 1\}} \Bigg\{ \mu_{ti}^m \theta_i 1_{m > N(\mathbf{1}\boldsymbol{\mu}_t, i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) - \sum_{l=m+1}^{N} \mu_{ti}^l \right) \theta_i 1_{m = N(\mathbf{1}\boldsymbol{\mu}_t, i)}$$

$$- \left[\hat{\mu}_{ti}^m \theta_i 1_{m > N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\hat{\boldsymbol{\mu}}_t + \mu_t^0) - \sum_{l=m+1}^{N} \hat{\mu}_{ti}^l \right) \theta_i 1_{m = N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} \right] \Bigg\}.$$

$$= \sum_{m \neq n, m=0}^{N} \sum_{i=1}^{I} \Bigg\{ \mu_{ti}^m \theta_i 1_{m > N(\mathbf{1}\boldsymbol{\mu}_t, i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) - \sum_{l=m+1}^{N} \mu_{ti}^l \right) \theta_i 1_{m = N(\mathbf{1}\boldsymbol{\mu}_t, i)}$$

$$- \left[\hat{\mu}_{ti}^m \theta_i 1_{m > N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\hat{\boldsymbol{\mu}}_t + \mu_t^0) - \sum_{l=m+1}^{N} \hat{\mu}_{ti}^l \right) \theta_i 1_{m = N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} \right] \Bigg\}.$$

$$= \overline{q} - \sum_{i=1}^{I} \left[\mu_{ti}^n \theta_i 1_{n > N(\mathbf{1}\boldsymbol{\mu}_t, i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\boldsymbol{\mu}_t + \mu_t^0) - \sum_{l=n+1}^{N} \mu_{ti}^l \right) \theta_i 1_{n = N(\mathbf{1}\boldsymbol{\mu}_t, i)} \right]$$

$$- \left\{ \overline{q} - \sum_{i=1}^{I} \left[\hat{\mu}_{ti}^n \theta_i 1_{n > N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} + \left(\alpha_i(\overline{q}, \mathbf{1}\hat{\boldsymbol{\mu}}_t + \mu_t^0) - \sum_{l=n+1}^{N} \hat{\mu}_{ti}^l \right) \theta_i 1_{n = N(\mathbf{1}\hat{\boldsymbol{\mu}}_t, i)} \right] \right\}.$$

We can bound this last expression by

$$\sum_{i=1}^{I}\left[\hat{\mu}_{ti}^{n}\theta_i 1_{n>N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)}+\left(\alpha_i(\overline{q},\mathbf{1}\hat{\boldsymbol{\mu}}_t+\mu_t^0)-\sum_{l=n+1}^{N}\hat{\mu}_{ti}^{l}\right)\theta_i 1_{n=N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)}\right].$$

Thus, ${}_uW^n(\boldsymbol{\mu}^n)-{}_uW^n(\hat{\boldsymbol{\mu}})\leq$

$$\begin{aligned}&\sum_{t=1}^{T}\beta^{t-T}p_t(\boldsymbol{\mu})\cdot\sum_{i=1}^{I}\left[\hat{\mu}_{ti}^{n}\theta_i 1_{n>N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)}+\left(\alpha_i(\overline{q},\mathbf{1}\hat{\boldsymbol{\mu}}_t+\mu_t^0)-\sum_{l=n+1}^{N}\hat{\mu}_{ti}^{l}\right)\theta_i 1_{n=N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)}\right]\\ \leq&\sum_{t=1}^{T}\beta^{t-T}\sum_{i=1}^{I}\left[\hat{\mu}_{ti}^{n}1_{n>N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)}+\left(\alpha_i(\overline{q},\mathbf{1}\hat{\boldsymbol{\mu}}_t+\mu_t^0)-\sum_{l=n+1}^{N}\hat{\mu}_{ti}^{l}\right)1_{n=N(\mathbf{1}\hat{\boldsymbol{\mu}}_t,i)}\right]\\ =&S(\hat{\boldsymbol{\mu}}),\end{aligned}$$

which then yields the first part of the lemma. If, in addition, Property II. is satisfied, then that expression equals

$$\sum_{i=1}^{I}\left(\hat{\mu}_{ti}^{n}-\mu_{ti}^{n}\right)\theta_i 1_{n>N(\mathbf{1}\boldsymbol{\mu}_t,i)},$$

which yields the second part of the lemma. □

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`