An ergodic theorem for mutual insurance

John Armstrong

Department of Mathematics, King's College London

Bachelier, Paris, 2026

Statement of result

Setup:

- Assume a discrete-time complete market with time points $\mathcal{T}$.
- Each agent has idiosyncratic risk given by their state. Independent of each other and the market.
- Each agent has an income stream determined by their state and the market's state.
- Each agent's preferences determined by their state, the market's state and their consumption, independent of other agents.
- At each time, n agents exchange derivatives/insurance contracts which are
  - **Admissible:** mutually self-financing
  - **Acceptable:** coalitionally stable, it is not in the interests of a subgroup of agents in a given state to reject the chosen contracts

Statement of result

Setup:

- Assume a discrete-time complete market with time points $\mathcal{T}$.
- Each agent has idiosyncratic risk given by their state. Independent of each other and the market.
- Each agent has an income stream determined by their state and the market's state.
- Each agent's preferences determined by their state, the market's state and their consumption, independent of other agents.
- At each time, n agents exchange derivatives/insurance contracts which are
  - **Admissible:** mutually self-financing
  - **Acceptable:** coalitionally stable, it is not in the interests of a subgroup of agents in a given state to reject the chosen contracts

Theorem

Under modest assumptions

- The value function for this problem tends to a unique limit as $n \to \infty$
- This limit can be achieved by a "simple strategy".

Simple strategies vs. real-world collective pensions

Simple strategy	Real-world collective pension
One-period Contracts only with agents in the same state, paying off at $t+1$.	**Multi-period** Payments now in return for payments at $t_{\mathrm{RA}}, t_{\mathrm{RA}}+1, \ldots$
Individual accounting Each agent's wealth is known at each time.	**Asset pooling** Only the fund's total is known; a member's share is up for debate.
Homogeneous Agents in the same state consume, invest and receive alike and only contract with agents in their state.	**Heterogeneous** Assets of different generations are managed together.

Simple strategies vs. real-world collective pensions

Simple strategy	Real-world collective pension
One-period Contracts only with agents in the same state, paying off at $t+1$.	**Multi-period** Payments now in return for payments at $t_{\mathrm{RA}}, t_{\mathrm{RA}}+1, \ldots$
Individual accounting Each agent's wealth is known at each time.	**Asset pooling** Only the fund's total is known; a member's share is up for debate.
Homogeneous Agents in the same state consume, invest and receive alike and only contract with agents in their state.	**Heterogeneous** Assets of different generations are managed together.

Simple strategies converge slowly: Insurance only occurs between agents in the same state.
Practical strategies: invest as you please, then at $t+1$ share the funds of the deceased in proportion to (wealth at $t+1$) $\times$ (odds of dying in $[t, t+1)$).

The mathematical challenges

The challenges:

- Dynamic programming is a backward-induction technique
- Agents' positions are independent at time 0, but if they exchange long-term contracts they will not be independent at time δt
- Establishing independence of each agent's strategy seems to need forward-induction.
- Coalitional stability is a non-convex condition, so we cannot use averaging arguments and Jensen's inequality to relate finite and infinite problems.

Proof strategy

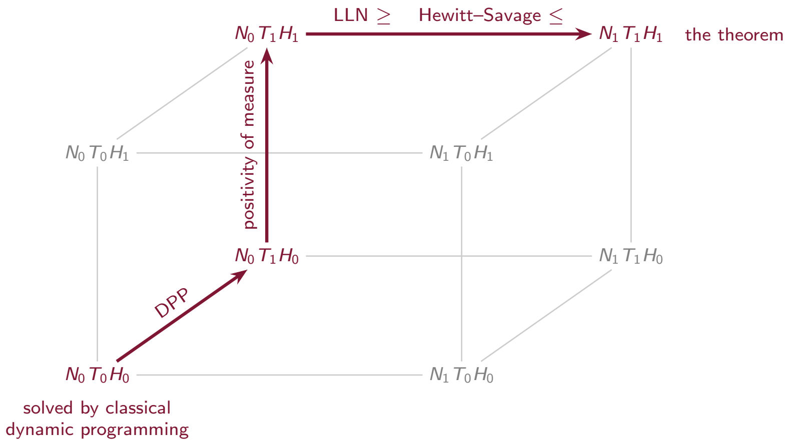

N_0/N_1: infinite/finite fund T_0/T_1: one-/multi-period H_0/H_1: homogeneous/heterogeneous

Agent preferences

Example

$\mathcal{J}(c) = \mathbb{E}_{\mathbb{P}}(\sum_{t\in\mathcal{T}} u(c_t))$

- Beliefs about market dynamics (e.g. choice of $\mathbb{P}$)
- Preferences over distribution (e.g. utility function u).

Example

$\mathcal{J}(c) = \mathbb{E}_{\mathbb{P}}(\sum_{t\in\mathcal{T}} u(c_t))$

- Beliefs about market dynamics (e.g. choice of $\mathbb{P}$)
- Preferences over distribution (e.g. utility function u).

No shared physical measure: Agents agree on prices $\mathbb{Q}^M$ but not on dynamics $\mathbb{P}^M$.

Agent preferences

Example

$\mathcal{J}(c) = \mathbb{E}_{\mathbb{P}}(\sum_{t \in \mathcal{T}} u(c_t))$

- Beliefs about market dynamics (e.g. choice of $\mathbb{P}$)
- Preferences over distribution (e.g. utility function u).

No shared physical measure: Agents agree on prices $\mathbb{Q}^M$ but not on dynamics $\mathbb{P}^M$.

No shared paradigm:

- Some agents may have recursive preferences
- Some agents may have S-shaped utilities
- Some agents may be robust optimizers

Agent preferences

Example

$\mathcal{J}(c) = \mathbb{E}_{\mathbb{P}}(\sum_{t \in \mathcal{T}} u(c_t))$

- Beliefs about market dynamics (e.g. choice of $\mathbb{P}$)
- Preferences over distribution (e.g. utility function u).

No shared physical measure: Agents agree on prices $\mathbb{Q}^M$ but not on dynamics $\mathbb{P}^M$.

No shared paradigm:

- Some agents may have recursive preferences
- Some agents may have S-shaped utilities
- Some agents may be robust optimizers

Axiomatic approach:

- $\mathcal{J}(c)$ is a value function acting on stochastic processes c representing income streams
- **Boring axioms:** Preferences over consumption are monotonic, non-satiating, appropriately continuous. . .
- **Invariance:** Agents don't care about each other's income (no altruism, no envy)
- **Compactness:** The individual investment problem is well-posed.

When are two probability spaces equivalent?

Mod 0 isomorphism: A measure preserving map $\phi : \Omega^1 \to \Omega^2$ with measurable inverse between two sets of full measure

Example

$([0,1], \lambda)$ and $([0,1] \times [0,1], \lambda)$ are mod 0 isomorphic.

Proof: define $\phi : [0,1] \times [0,1] \to [0,1]$ *by sending* (x, y) *to the point given by alternating the binary digits of* x *and* y*.*

When are two probability spaces equivalent?

Mod 0 isomorphism: A measure preserving map $\phi : \Omega^1 \to \Omega^2$ with measurable inverse between two sets of full measure

Example

$([0,1], \lambda)$ and $([0,1] \times [0,1], \lambda)$ are mod 0 isomorphic.

Proof: define $\phi : [0,1] \times [0,1] \to [0,1]$ *by sending* (x, y) *to the point given by alternating the binary digits of* x *and* y*.*

Stable classification of probability spaces: If Ω^1 and Ω^2 are probability spaces
$\Omega^1 \times [0,1] \cong \Omega^2 \times [0,1]$.
The stable classification ignores atoms.

When are two probability spaces equivalent?

Mod 0 isomorphism: A measure preserving map $\phi : \Omega^1 \to \Omega^2$ with measurable inverse between two sets of full measure

Example

$([0,1], \lambda)$ and $([0,1] \times [0,1], \lambda)$ are mod 0 isomorphic.

Proof: define $\phi : [0,1] \times [0,1] \to [0,1]$ *by sending* (x, y) *to the point given by alternating the binary digits of* x *and* y*.*

Stable classification of probability spaces: If Ω^1 and Ω^2 are probability spaces
$\Omega^1 \times [0,1] \cong \Omega^2 \times [0,1]$.
The stable classification ignores atoms.

Stable classification of filtered probability spaces:

- Let $\mathcal{T}$ be a set of time points.
- Let $\Omega^{\text{can}}_{\mathcal{T}} := [0,1]^{\mathcal{T}}$
- Any filtered probability space indexed by $\mathcal{T}$ is stably equivalent to $\Omega^{\text{can}}_{\mathcal{T}} := [0,1]^{\mathcal{T}}$.

Definition of invariant preferences

- Let Ω^1 be a filtered probability space containing "relevant data", e.g. the state of the market and the agent under consideration.
- *Invariant preferences* for Ω^1 are given by a value function $\mathcal{J}$ on stochatic processes defined on $\Omega^1 \times \Omega^{\text{can}}_{\mathcal{T}}$ which is invariant under mod 0 isomorphisms of the factor on the right.
- Given Ω^2 a filtered probability space containing "irrelevant data", e.g. the state of other agents define induced $\mathcal{J}'$ for processes c on $\Omega^1 \times \Omega^2$ by
  - Extend the processes trivially to processes $\bar{c}$ on $\Omega^1 \times \Omega^2 \times \Omega^{\text{can}}_{\mathcal{T}}$
  - Use the isomorphism $\Omega^1 \times \Omega^2 \times \Omega^{\text{can}}_{\mathcal{T}} \cong \Omega^1 \times \Omega^{\text{can}}_{\mathcal{T}}$ to obtain a process $\hat{c}$ defined on $\Omega^1 \times \Omega^{\text{can}}_{\mathcal{T}}$.
  - Define $\mathcal{J}'(c) = \mathcal{J}(\hat{c})$.

The Hewitt–Savage Zero-One Law

- Let $(X^i)_{i=1}^{\infty}$ be independent, identically distributed random variables.
- Let A be an element of the sigma algebra generated by the X^i.
- We say A is exchangeable if its value is invariant under finite permutations of the X^i.

Theorem

If A is exchangeable, $\mathbb{P}(A) = 1$ or $\mathbb{P}(A) = 0$.

The Hewitt–Savage Zero-One Law

- Let $(X^i)_{i=1}^{\infty}$ be independent, identically distributed random variables.
- Let A be an element of the sigma algebra generated by the X^i.
- We say A is exchangeable if its value is invariant under finite permutations of the X^i.

Theorem

If A is exchangeable, $\mathbb{P}(A) = 1$ or $\mathbb{P}(A) = 0$.

Corollary

Any random variable which is measurable in the *exchangeable sigma-algebra* is constant.

The Hewitt–Savage Zero-One Law

- Let $(X^i)_{i=1}^{\infty}$ be independent, identically distributed random variables.
- Let A be an element of the sigma algebra generated by the X^i.
- We say A is exchangeable if its value is invariant under finite permutations of the X^i.

Theorem

If A is exchangeable, $\mathbb{P}(A) = 1$ or $\mathbb{P}(A) = 0$.

Corollary

Any random variable which is measurable in the *exchangeable sigma-algebra* is constant.

Proof sketch that $N_1 T_1 H_1 \leq N_0 T_1 H_1$

- Use the symmetry of our problem to construct highly symmetric solutions for each n.
- Use compactness to show these solutions tend to a limit as $n \to \infty$.
- Observe that these solutions are measurable in an exchangeable sigma-algebra
- Use a conditional Hewitt–Savage Theorem to show that the limiting strategy lies in $N_0 T_1 H_1$.

Intuition: the collapse of degrees of freedom

- An infinite number of independent degrees of freedom is incompatible with compactness.
- For compact problems, as $n \to \infty$ there must be a collapse in the number of degrees of freedom.
- The freedom to choose heterogeneous strategies vanishes, all strategies become homogeneous.
- In thermodynamics: a gas trapped in a box becomes homogeneous.

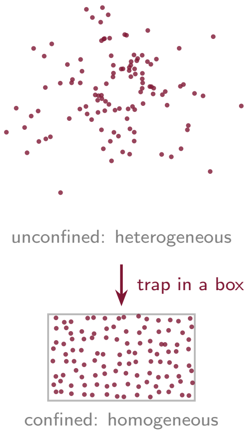

The probability spaces for our problems

- Let Ω^M represent the market. Let Ω^i represent the state of individual i.
- Let $\Omega^{\text{can}}_{\mathcal{T}}$ be a generic probability space containing irrelevant information.
- $N^1T^1H^1$: A set of contracts $\gamma^i_{s,t}$ is a set of random variables defined on

$$(\Omega^M \times \Omega^{\text{can}}_{\mathcal{T}}) \times \prod_{i=1}^{n} (\Omega^i \times \Omega^{\text{can}}_{\mathcal{T}})$$

  indexed by individual i, time the contract is agreed s, time the payment will be received t.
- $N^0T^0H^0$: Investment and consumption streams defined on

$$\Omega^M \times \Omega^1 \times \Omega^{\text{can}}_{\mathcal{T}}$$

- Symmetrise, apply Hewitt–Savage and classification

$$\Omega^{\text{can}}_{\mathcal{T}} \times (\Omega^M \times \Omega^{\text{can}}_{\mathcal{T}}) \times \Omega^1 \times \Omega^{\text{can}}_{\mathcal{T}} \times \prod_{i=2}^{\infty} (\Omega^i \times \Omega^{\text{can}}_{\mathcal{T}})$$

Compactness - the analytic details needed

- Compactness is a property of preferences: roughly speaking it requires that if c_n is a sequence of income streams where $\mathcal{J}(c_n)$ is bounded below and $\mathbb{E}_{\mathbb{Q}}(c_n)$ is bounded above then c_n has a convergent subsequence.
- Compactness is effectively well-posedness of the easy problem $N_0 H_0 T_0$.

Compactness in this sense is formally defined in the paper and is the key technical assumption.

- Establishing compactness is challenging: Kramkov and Schachermayer's work establishes compactness for von Neumann–Morgenstern functions in one-period models subject to certain moment estimates.
- Paper extend this to various recursive preferences of interest in pensions.

Summary

- The optimal pension design subject to coalitional stability is given by a "DC + longevity insurance" design.
- It is impossible to share market risks between generations in complete markets, contradicting the draft report of the UK pension commission.
- The proof depends upon symmetry arguments and compactness.

https://arxiv.org/abs/2608.14256