Recursive Preferences, Correlation Aversion, and the Temporal Resolution of Uncertainty

Lorenzo Stanca

ESOMAS Department and Collegio Carlo Alberto
Università di Torino

Teaching slides

Motivation

Many economic decisions involve **risks that persist over time**:

- Long-run income risk and human capital.
- Persistent productivity or growth shocks in asset pricing.
- Climate risks.

Standard models (expected utility with exponential discounting):

- mix together **risk aversion** and **intertemporal substitution**;
- have little to say about attitudes toward **persistence / correlation** of risks.

Recursive preferences do not have these limitations.

What the paper does

1. Introduces a clear behavioral notion of **correlation aversion**: dislike of consumption risks that are positively correlated over time.

2. Shows that, for a broad class of recursive preferences:
    - correlation aversion $\Longleftrightarrow$ **increasing relative risk aversion (IRRA)**.
    - Surprisingly, the paper shows that correlation aversion is connected to model uncertainty/robustness.

3. Introduces a measure of correlation aversion: the **persistence premium**: How much consumption would you sacrifice to remove persistence?
    - Provides an Arrow–Pratt style approximation.
    - The approximation shows how correlation aversion depends on the relevant parameters.

4. Applications:
    - Asset pricing/portfolio choice and the equity premium.
    - Optimal income taxation and social mobility.

A simple example

Two consumption lotteries, same marginal distribution, different correlation:

Gamble A: "once and for all"

- At $t = 1$ a fair coin is tossed:

$$\text{Heads} \Rightarrow (c_1, c_2, c_3, \dots) = (1, 1, 1, \dots)$$

$$\text{Tails} \Rightarrow (c_1, c_2, c_3, \dots) = (0, 0, 0, \dots)$$

- All uncertainty is resolved at $t = 1$. Risk is **fully persistent**.

Gamble B: i.i.d. each period

- At each $t = 1, 2, 3, \dots$ a fair coin is tossed: $c_t = 1$ or $c_t = 0$ with probability $1/2$.
- **No persistence**: consumption each period is independent.

Question: Which path do we prefer?

Two motives in conflict

1. **Hedging motive**:
   - In B, good and bad outcomes are spread over time.
   - In A, you are either always rich or always poor.
   - A risk-averse agent (with concave utility) prefers B for hedging.

2. **Information / timing motive**:
   - In A, you quickly learn your entire consumption path.
   - In B, uncertainty is resolved gradually.
   - Recursive preferences often feature a **preference for early resolution**: the agent likes to know sooner.

3. So:
   - A is **better for information**, worse for hedging.
   - B is **better for hedging**, worse for information.

4. **Correlation aversion**: the hedging motive dominates $\Rightarrow$ prefer B to A.

Recursive preferences with finite horizon

- Time is discrete: $t = 0, 1, \ldots, T$.
- Consumption at t is c_t.
- **Kreps–Porteus (KP) recursive utility**:

$$V_T(c_T) = u(c_T),$$

$$V_t(c_t, m_{t+1}) = u(c_t) + \beta\, \phi^{-1}\big(\mathbb{E}_{m_{t+1}}[\phi(V_{t+1})]\big).$$

- Components:
    - u: period utility, governs **intertemporal substitution** (EIS).
    - ϕ: **risk adjustment** over future utilities, governs risk attitude.
    - β: time discount factor.
- Key feature: separates **risk aversion** (via ϕ) from **EIS** (via u).

A familiar special case: Epstein–Zin

Epstein–Zin preferences fit into KP with

$$u(c) = \frac{c^\rho}{\rho}, \quad \phi(x) = \frac{(\rho x)^{\alpha/\rho}}{\alpha},$$

where

- $1 - \alpha$ is **relative risk aversion**,
- $\frac{1}{1-\rho}$ is the **elasticity of intertemporal substitution (EIS)**.

Widely used in:

- asset pricing (long-run risk, rare disasters),
- macro and fiscal policy,
- growth and climate models.

This paper shows that, in such models, the shape of ϕ is crucial for **correlation aversion**.

Early resolution of uncertainty

An information structure is **more informative** in the Blackwell sense if it gives a finer signal about future consumption.

Preference for early resolution (PERU):

- If two consumption paths have the same distribution, but one reveals the future earlier in the Blackwell sense, the agent prefers it.
- Recursive preferences can exhibit PERU. This preference is linked to the curvature of ϕ.

Measuring PERU

The paper defines a local measure of preference for PERU (call it ER_ϕ) derived from the second derivative of ϕ:

$$ER_\phi \approx -\frac{\phi''(\text{today's utility})}{\phi'(\text{today's utility})} + \beta \frac{\phi''(\text{future utility})}{\phi'(\text{future utility})}.$$

Intuition: The Arrow-Pratt index for the present exceeds the discounted Arrow-Pratt measure for the future:

- If ER_ϕ is large and positive, the agent values early resolution a lot.

Increasing correlation over time: IECIT

Step 1: start from iid. Fix the one-period distribution ℓ and consider the iid process $d^{iid}(\ell)$: each period has distribution ℓ and today does not affect tomorrow.

Step 2: add correlation without changing one-period risk. An IECIT is a small perturbation of the joint law of (c_t, c_{t+1}) that:

- makes "staying the same" more likely $(c_{t+1} = c_t)$,
- makes "switching" less likely (high→low, low→high),
- <u>while keeping the marginal distribution in each period equal to ℓ.</u>

Interpretation. Same risk each period, but more persistence: good times (bad times) are more likely to be followed by good times (bad times).

Correlation order and correlation aversion

- **Correlation order.** Write $d \geq_C d'$ if d can be obtained from d' by a finite sequence of IECIT steps. Interpretation: d is (weakly) more persistent / more positively correlated over time than d'.

- **Correlation aversion (formal axiom).** Preferences are correlation averse if, whenever we start from the iid benchmark and increase persistence (in the sense of $\geq_C$), the decision maker is weakly worse off:

$$d \geq_C d' \geq_C d^{iid}(\ell) \quad \Longrightarrow \quad d^{iid}(\ell) \ \succeq \ d' \ \succeq \ d.$$

Main theoretical result

Key theorem (informal)

Within KP recursive preferences, suppose the agent always prefers early resolution of uncertainty (PERU). Then:

Correlation aversion

$$\iff$$

ϕ is concave and has **increasing relative risk aversion (IRRA)**.

- Under IRRA, I show that ER_ϕ is **restricted**: you cannot value early resolution "too much" given your level of risk aversion.
- Epstein–Zin and Hansen–Sargent multiplier preferences satisfy these conditions.
- A linear ϕ (no risk aversion over future utility) $\Rightarrow$ **no correlation aversion**.

Connection with Model uncertainty / robustness

The paper further shows:

- a strengthened version of IRRA (called strong correlation aversion) implies a representation with **model misspecification concerns**.

Recursive utility can be written as:

$$V_t(c_t, m_{t+1}) = u(c_t) + \beta \min_{\ell} \left\{ \mathbb{E}_{\ell}[V_{t+1}] + I_t(\ell \| m_{t+1}) \right\},$$

where I_t is a statistical distance ("cost") between models ℓ and m_{t+1}.

Interpretation:

- correlation aversion is connected to **robustness** to misspecification of the consumption process.

Persistence premium: what is it?

Thought experiment

- Compare two consumption processes:
  1. A **correlated** process with persistence parameter $\varepsilon \in [0,1]$: current shocks partly carry over to the future.
  2. An **i.i.d.** process with the same marginal distribution each period.
- To make the two indifferent, scale down consumption in the i.i.d. process by a factor $(1-\pi(\varepsilon))$.

- $\pi(\varepsilon) =$ **persistence premium**:

  fraction of consumption you are willing to give up
  to eliminate persistence (correlation) in consumption.
- $\pi(\varepsilon)$ increases with
  - the degree of persistence ε,
  - risk aversion (curvature of ϕ),
  - and the variance of consumption.

Approximate formula (Epstein–Zin case)

In an Epstein–Zin specification with:

- relative risk aversion $1-\alpha$,
- EIS $\frac{1}{1-\rho}$,
- two consumption levels c_H and c_L,
- persistence parameter ε,

the paper derives an approximation:

$$\pi(\varepsilon) \approx \tilde{a} + \tilde{b}\,\varepsilon\Big(1-\frac{\alpha}{\rho}\Big)\Big(\frac{1}{c_H}+\frac{1}{c_L}\Big) - \tilde{c}\,(\varepsilon^2-1)\,ER_\phi,$$

where $\tilde{a}$, $\tilde{b}$, $\tilde{c} > 0$ are constants.

Takeaways:

- More persistence (ε) $\Rightarrow$ higher persistence premium.
- Higher risk aversion $\Rightarrow$ premium rises faster.
- Stronger preference for information (high ER_ϕ) moderates the increase.

Application I: asset pricing/portfolio choice

Empirical facts:

- Stock returns are procyclical and volatile.
- The **equity premium** is high relative to standard models.

Long-run risk models (e.g. Bansal–Yaron):

- consumption growth has a small but **persistent** component,
- representative agent has Epstein–Zin preferences.

With correlation aversion:

- persistent long-run risk makes equities particularly bad hedges,
- investors demand a higher equity premium as compensation.

In this case, the persistence premium quantifies how costly this long-run risk is in utility terms.

How large is the persistence premium?

Calibrate a standard long-run risk model with Epstein–Zin preferences.

Compute the persistence premium:

$$\pi = 1 - \frac{U_0(\text{correlated LRR process})}{U_0(\text{i.i.d. version})}.$$

For parameter values commonly used in the literature, the paper finds:

- π can be around 30–40%.

Interpretation:

- The representative agent would give up 30–40% of lifetime consumption to shut down the persistent component of consumption growth.
- This looks implausibly large from an introspective or experimental viewpoint.

Tension with experimental evidence

Experiments (e.g. Rohde–Yu) can directly elicit **correlation aversion**.

Findings:

- correlation aversion exists but is **moderate**,
- compatible with relatively modest risk aversion (say, between 1 and 2).

Why the discrepancy?

- In standard Epstein–Zin, increasing risk aversion also increases the preference for early resolution of uncertainty.
- To generate enough correlation aversion, you end up with unrealistically high risk aversion and enormous persistence premia.

A more flexible risk aggregator (HARA)

The paper generalizes ϕ beyond Epstein–Zin to a **HARA** (hyperbolic absolute risk aversion) form:

$$\phi_{\gamma,b}(x) = \frac{1-\gamma}{\gamma}\left(\frac{x}{1-\gamma} + b\right)^{\gamma}, \quad 0 \neq \gamma < 1,\ b \geq 0.$$

Epstein–Zin corresponds to $b = 0$ and a particular γ.

Advantages:

- risk aversion and preference for early resolution can be adjusted more independently;
- can match experimentally observed correlation aversion with **lower and more realistic** risk aversion.

Hence, you can rationalize bond vs. stock holdings without relying on high risk aversion.

Application II: income taxation and social mobility

Start from a dynamic model of optimal progressive income taxation (Benabou; Park):

- dynasties with human capital h_t, income, labor choice, education choice;
- human capital may be persistent across generations;
- government chooses a progressive tax rule characterized by a parameter τ.

Under **standard expected utility**:

- the optimal degree of progressivity is roughly constant when human capital becomes more persistent.

With correlation-averse **recursive preferences**:

- greater persistence of ability/human capital $\Rightarrow$ much higher optimal progressivity;
- progressive taxes act like an **inheritance tax** conditional on the past history of human capital of the individual's family.

Correlation aversion and social mobility

Persistent human capital means:

- high-human-capital families stay rich across generations,
- low-human-capital families stay poor.

Correlation-averse preferences dislike this **persistent inequality** in lifetime consumption paths.

The planner therefore:

- chooses higher τ (more progressive taxes) when persistence is high;
- effectively reduces correlation between parental and children outcomes;
- increases **social mobility**.

Policy message:

- with correlation aversion, it is optimal to target **long-run inequalities** (education, inheritance, mobility), not only short-run income shocks. This is not true with standard utility!

Conclusion

1. **Correlation aversion** is a distinct and economically important attitude: dislike of persistent, highly correlated consumption risks.

2. For a broad class of recursive preferences that like early resolution, correlation aversion $\Longleftrightarrow$ **IRRA** of the risk aggregator ϕ.

3. The **persistence premium** provides a quantitative measure of correlation aversion, highlighting the trade-off between hedging and non-instrumental information.

4. Correlation aversion has strong implications for:
    - asset pricing (equity premium, long-run risk),
    - optimal taxation and social mobility,
    - models with robustness to misspecification.

Open questions / further research

- **Measurement / identification:** can we empirically separate correlation aversion from (i) standard risk aversion and (ii) preferences for early resolution?
- **Disciplining macro-finance calibrations:** can long-run risk / disaster models match equity premium while keeping the implied persistence premium in a plausible range?
- **Theory extensions:** extend the IECIT / correlation order (continuous time or continuum of shocks) and characterize which recursive aggregators preserve monotonicity in that order.