Decision Theory and Stochastic Growth

Arthur Robson, Larry Samuelson, Jakub Steiner

Online, October 2022

Stochastic Growth

macroeconomic behavior $\Leftrightarrow$ individual information processing

Representations with "behavioral flavor"

In the context of stochastic growth, long strikes of luck

- are exponentially unlikely
- but provide an exponential wealth advantage

$\Rightarrow$ Large deviations from law of large numbers

Equivalent tradeoffs:

- growth advantage vs rarity of a deviation
- benefit vs cost of information

Our Results

growth process	**known representation**
idiosyncratic risk	expected-utility maximization
aggregate risk	growth-optimal portfolio

Our Results

growth process	known representation	new representation
idiosyncratic risk	expected-utility maximization	wishful thinking
aggregate risk	growth-optimal portfolio	rational inattention

These findings

- shed light on decision models of interest
- shed light on the growth process
- transfer solution methods across domains

Literature

Growth:
growth-optimal portfolios: Kelly'56,...
evolution: Blume&Easly'92, Robson'96, Robson&Samuelson'11,...

Biology: *hedging is learning*: Kussell&Leibler'05

Decision Theory:
wishful thinking: Caplin&Leahy'19
robust control: Hansen&Sargent'08
rational inattention: Caplin&Dean&Leahy'22, Matějka&McKay'15

Large Deviation Theory: Donsker&Varadhan Lemma, Sanov Theorem

1. Equivalences
2. Large Deviation Theory
   - tutorial
3. Choice, Uncertainty and Growth
   - idiosyncratic risk
   - aggregate risk
4. Applications
   - side information
   - misspecification

Wishful-Thinking

psychological model by Caplin&Leahy'19

Primitives: $a \in A$, $u(a, \boldsymbol{\theta})$, $\boldsymbol{\theta} \sim p \in \Delta(\Theta)$

Wishful thinker:

- chooses a subjective belief q at distortion cost $D_{KL}(q \parallel p)$
- chooses action a
- enjoys subjective expectation $\mathrm{E}_q\, u(a, \boldsymbol{\theta})$

$$\max_{a,q} \left\{ \mathrm{E}_q\, u(a, \boldsymbol{\theta}) - D(q \parallel p) \right\}$$

Proposition

Wishful thinker chooses action a^*

$\Leftrightarrow$

EU maximizer with prior p and utility $U(a, \theta) := e^{u(a,\theta)}$ chooses a^*.

See Strzalecki'11 for a related result

Growth-Optimal Portfolio

investment model of Kelly'56

Investor allocates fractions $\alpha(a)$ to assets $a \in A$; $\alpha \in \Delta(A)$

Asset a has stochastic return $e^{u(a,\boldsymbol{\theta})}$; $\boldsymbol{\theta} \sim p \in \Delta(\Theta)$

- $u(a, \boldsymbol{\theta})$ is the growth rate of asset a

Portfolio α enjoys return $\mathrm{E}_\alpha\, e^{u(\mathbf{a},\theta)}$ in state θ

Kelly: investor should maximize expected growth rate

$$\max_{\alpha} \mathrm{E}_p \ln \mathrm{E}_\alpha\, e^{u(\mathbf{a},\boldsymbol{\theta})}$$

Growth-optimal portfolio a.s. outgrows other portfolios for iid $\boldsymbol{\theta}_t \sim p$

Growth-Optimal Portfolio

investment model of Kelly'56

Investor allocates fractions $\alpha(a)$ to assets $a \in A$; $\alpha \in \Delta(A)$

Asset a has stochastic return $e^{u(a,\boldsymbol{\theta})}$; $\boldsymbol{\theta} \sim p \in \Delta(\Theta)$

- $u(a, \boldsymbol{\theta})$ is the growth rate of asset a

Portfolio α enjoys growth rate $\ln \mathrm{E}_\alpha \, e^{u(\mathbf{a},\theta)}$ in state θ

Kelly: investor should maximize expected growth rate

$$\max_\alpha \mathrm{E}_p \ln \mathrm{E}_\alpha \, e^{u(\mathbf{a},\boldsymbol{\theta})}$$

Growth-optimal portfolio a.s. outgrows other portfolios for iid $\boldsymbol{\theta}_t \sim p$

Rational-Inattention Problem
information-acquisition model of Matějka&McKay'15

Trade-off between benefit and cost of information

Optimization over state-dependent stochastic choice rules $q(a \mid \theta)$

$$\max_{q \in \Delta(A)^{\Theta}} \left\{ \mathrm{E}_{p,q}\, u(\mathbf{a}, \boldsymbol{\theta}) - I_{p,q}(\boldsymbol{\theta}; \mathbf{a}) \right\}$$

Let $q(a) := \mathrm{E}_p\, q(a \mid \boldsymbol{\theta})$ be the marginal optimal choice rule

Proposition

A mixed strategy is a growth-optimal portfolio
$\Leftrightarrow$
it is the marginal optimal choice rule in the rational-inattention problem.

Donsker-Varadhan Lemma

$$\ln \mathsf{E}_p\, e^{u(\mathbf{x})} = \max_{q \in \Delta(X)} \left\{ \mathsf{E}_q\, u(\mathbf{x}) - D(q \parallel p) \right\}$$

growth rate of expectation

=

expectation of growth rate at an optimized biased belief less of a penalty

Large Deviation Principle

T iid draws $\mathbf{x}_t \sim p \in \Delta(X)$,

Empirical distribution $q_T(x) = \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}_{\mathbf{x}_t = x}$

LLN: it is likely that $q_T \approx p$
The event $q_T(x) = q \neq p$ is called large deviation

Large Deviation Principle $\Pr(q_T = q) \approx e^{-T \times I(q)}$ for a rate function $I(q)$

Sanov Theorem: $I(q) = D(q \parallel p)$

Implication for Stochastic Growth

- consider stochastic growth rate $u(\mathbf{x}_t)$ in each period t
- $\mathbf{w}_T = \prod_{t=1}^{T} e^{u(\mathbf{x}_t)}$
- $\mathrm{E}\,\mathbf{w}_T = \left(\mathrm{E}_p\, e^{u(\mathbf{x})}\right)^T$

Varadhan's Lemma: $\ln \mathrm{E}_p\, e^{u(\mathbf{x})} = \max_q \left\{ \mathrm{E}_q\, u(\mathbf{x}) - I(q) \right\}$

- sum over all sequences $x^T = (x_1, \ldots, x_T)$
- $w_T(x^T) = \exp\left[T \times \mathrm{E}_{q_T}\, u(\mathbf{x})\right]$
- probability of empirical distribution q is $\approx \exp\left[-T \times I(q)\right]$
- $\mathrm{E}\,\mathbf{w}_T \approx \int_{q \in \Delta(X)} \exp\left[T\left(\mathrm{E}_q\, u(\mathbf{x}) - I(q)\right)\right] dq$
- $\mathrm{E}\,\mathbf{w}_T \approx \exp\left[T \max_q \left\{\mathrm{E}_q\, u(\mathbf{x}) - I(q)\right\}\right]$

Implication for Stochastic Growth

- consider stochastic growth rate $u(\mathbf{x}_t)$ in each period t
- $\mathbf{w}_T = \prod_{t=1}^{T} e^{u(\mathbf{x}_t)}$
- growth rate $\ln \mathrm{E}_p\, e^{u(\mathbf{x})}$

Varadhan's Lemma: $\ln \mathrm{E}_p\, e^{u(\mathbf{x})} = \max_q \left\{ \mathrm{E}_q\, u(\mathbf{x}) - I(q) \right\}$

- sum over all sequences $x^T = (x_1, \ldots, x_T)$
- $w_T(x^T) = \exp\left[T \times \mathrm{E}_{q_T}\, u(\mathbf{x})\right]$
- probability of empirical distribution q is $\approx \exp\left[-T \times I(q)\right]$
- $\mathrm{E}\,\mathbf{w}_T \approx \int_{q \in \Delta(X)} \exp\left[T\left(\mathrm{E}_q\, u(\mathbf{x}) - I(q)\right)\right] dq$
- $\mathrm{E}\,\mathbf{w}_T \approx \exp\left[T \max_q \left\{\mathrm{E}_q\, u(\mathbf{x}) - I(q)\right\}\right]$

Implication for Stochastic Growth

- consider stochastic growth rate $u(\mathbf{x}_t)$ in each period t
- $\mathbf{w}_T = \prod_{t=1}^{T} e^{u(\mathbf{x}_t)}$
- growth rate $\ln \mathrm{E}_p\, e^{u(\mathbf{x})}$

Varadhan's Lemma: $\ln \mathrm{E}_p\, e^{u(\mathbf{x})} = \max_q \left\{ \mathrm{E}_q\, u(\mathbf{x}) - I(q) \right\}$

- sum over all sequences $x^T = (x_1, \ldots, x_T)$
- $w_T(x^T) = \exp\left[T \times \mathrm{E}_{q_T}\, u(\mathbf{x})\right]$
- probability of empirical distribution q is $\approx \exp\left[-T \times I(q)\right]$
- $\mathrm{E}\,\mathbf{w}_T \approx \int_{q \in \Delta(X)} \exp\left[T\left(\mathrm{E}_q\, u(\mathbf{x}) - I(q)\right)\right] dq$
- $\mathrm{E}\,\mathbf{w}_T \approx \exp\left[T \max_q \left\{\mathrm{E}_q\, u(\mathbf{x}) - I(q)\right\}\right]$

Implication for Stochastic Growth

Combine Sanov and Varadhan

Donsker-Varadhan formula:

$$\ln \mathrm{E}_p \, e^{u(\mathbf{x})} = \max_{q \in \Delta(X)} \left\{ \mathrm{E}_q \, u(\mathbf{x}) - D(q \parallel p) \right\}$$

"All" growth is supported by a single large deviation

growth advantage vs likelihood of the deviation

Setting

Continuous population, discrete time

Each t, each agent chooses action a, her wealth multiplies by $e^{u(a,\boldsymbol{\theta}_t)}$

iid $\boldsymbol{\theta}_t \sim p \in \Delta(\Theta)$

- idiosyncratic risk: each agent is subject to her own draw
- aggregate risk: each agent is subject to a same draw

We maximize growth rate of the aggregate wealth

Optimal strategies are

- visible
- likely to be reproduced

Two Problems

Idiosyncratic risk: every shock sequence is experienced by a subpopulation

$$\max_{a \in A} \ln \mathrm{E}_p \, e^{u(a,\boldsymbol{\theta})}$$

$\Rightarrow$ EUT with $U = e^u$

Aggregate risk: whole population experiences a same shock sequence

$$\max_{\alpha \in \Delta(A)} \mathrm{E}_p \ln \mathrm{E}_\alpha \, e^{u(\mathbf{a},\boldsymbol{\theta})}$$

$\Rightarrow$ Kelly's hedging on population level

Two Problems

Idiosyncratic risk: every shock sequence is experienced by a subpopulation

$$\max_{a \in A} \mathrm{E}_p \, e^{u(a,\boldsymbol{\theta})}$$

$\Rightarrow$ EUT with $U = e^u$

Aggregate risk: whole population experiences a same shock sequence

$$\max_{\alpha \in \Delta(A)} \mathrm{E}_p \ln \mathrm{E}_\alpha \, e^{u(\mathbf{a},\boldsymbol{\theta})}$$

$\Rightarrow$ Kelly's hedging on population level

Wishful Thinking $\Leftrightarrow$ EUT

proof

$$\arg\max_a \mathrm{E}_p\, e^{u(a,\boldsymbol{\theta})}$$

Donsker-Varadhan transforms this into the wishful-thinking problem

$$\arg\max_{a,q} \{\mathrm{E}_q\, u(a,\boldsymbol{\theta}) - D(q \parallel p)\}$$

a deviation q is enjoyed by fraction $e^{-T\times D(q\|p)}$ of population

wealth of the "lucky" subpopulation multiplies by $e^{T\times \mathrm{E}_q\, u(a,\boldsymbol{\theta})}$

An individual owns nontrivial wealth share only if

- is lucky to enjoy the optimal deviation
- has chosen optimal action against this deviation

Wishful Thinking $\Leftrightarrow$ EUT

proof

$$\arg\max_a \ln \mathrm{E}_p\, e^{u(a,\boldsymbol{\theta})}$$

Donsker-Varadhan transforms this into the wishful-thinking problem

$$\arg\max_{a,q} \{\mathrm{E}_q\, u(a,\boldsymbol{\theta}) - D(q \parallel p)\}$$

a deviation q is enjoyed by fraction $e^{-T \times D(q \| p)}$ of population

wealth of the "lucky" subpopulation multiplies by $e^{T \times \mathrm{E}_q\, u(a,\boldsymbol{\theta})}$

An individual owns nontrivial wealth share only if

- is lucky to enjoy the optimal deviation
- has chosen optimal action against this deviation

Sampled Distribution

Sample a $\$$ uniformly at a period t. The dollar's owner's shock $\theta_t \sim$

$$\tilde{q}(\theta) = \frac{p(\theta)e^{u(a^*,\theta)}}{\mathsf{E}_p\, e^{u(a^*,\boldsymbol{\theta})}}$$

Corollary

Sampled distribution $\tilde{q} =$ distribution q^* chosen by the wishful thinker.

Wealth Concentration

Retrospective growth rate $\mathrm{E}_{q^*}\, u(a^*, \boldsymbol{\theta})$

Growth rate of aggregate wealth is only $\mathrm{E}_{q^*}\, u(a^*, \boldsymbol{\theta}) - D(q^* \parallel p)$

$\Rightarrow$ Exponentially increasing concentration of wealth

Fraction $e^{-t \times D(q^* \| p)}$ of population owns “all” wealth in period t

expense of wishful thinking $=$ rate of wealth concentration

Robust Control

q^* – the sampled distribution

a^* – the chosen action

Proposition

$$p \in \arg\min_{p'} \left\{ \mathrm{E}_{p'}\, u(a^*, \boldsymbol{\theta}) + D\left(p' \parallel q^*\right) \right\}$$

Akin to the robust-control approach by Hansen&Sargent:

- agent doesn't quite trust q^*
- adopts the worst belief "nearby" q^*
- relaxed problem with shadow price normalized to 1

Growth-optimal Portfolio $\Leftrightarrow$ Rational Inattention
proof

Growth-optimal portfolio

$$\max_{\alpha} \mathrm{E}_p \ln \mathrm{E}_\alpha \, e^{u(\mathbf{a},\boldsymbol{\theta})}$$

Growth-optimal Portfolio ⇔ Rational Inattention

proof

Apply Donsker-Varadhan to each state θ

$$\max_{\alpha, \boldsymbol{q}} \left\{ \mathrm{E}_{p,q}\, u(\mathbf{a}, \boldsymbol{\theta}) - \mathrm{E}_p\, D\big(q(a \mid \boldsymbol{\theta}) \parallel \alpha(a)\big) \right\}$$

Growth-optimal Portfolio $\Leftrightarrow$ Rational Inattention

proof

Rewrite as the divergence of joint distributions

$$\max_{\alpha,\boldsymbol{q}} \left\{ \mathsf{E}_{\boldsymbol{p},\boldsymbol{q}}\, u(\mathbf{a},\boldsymbol{\theta}) - D\big(q(\theta,a) \parallel p(\theta)\alpha(a)\big) \right\}$$

Growth-optimal Portfolio ⇔ Rational Inattention

proof

Separate the divergence of the marginal action distributions

$$\max_{\alpha, q} \left\{ \mathsf{E}_{p,q}\, u(\mathbf{a}, \boldsymbol{\theta}) - \mathsf{E}_{q(a)}\, D\big(q(\theta \mid \mathbf{a}) \parallel p(\theta)\big) - D\big(q(a) \parallel \alpha(a)\big) \right\}$$

Marginal action distributions must coincide

$$\max_{\alpha,\boldsymbol{q}} \left\{ \mathsf{E}_{p,q}\, u(\mathbf{a},\boldsymbol{\theta}) - \mathsf{E}_{q(a)}\, D\big(q(\theta \mid \mathbf{a}) \parallel p(\theta)\big) - D\big(q(a) \parallel \alpha(a)\big) \right\}$$

Growth-optimal Portfolio $\Leftrightarrow$ Rational Inattention
proof

Drop optimization over α

$$\max_{q} \left\{ \mathsf{E}_{p,q}\, u(\mathbf{a}, \boldsymbol{\theta}) - \mathsf{E}_{q(a)}\, D\big(q(\theta \mid \mathbf{a}) \parallel p(\theta)\big) \right\}$$

Growth-optimal Portfolio ⇔ Rational Inattention
proof

This is the rational-inattention problem, as needed

$$\max_q \{\mathrm{E}_{p,q}\, u(\mathbf{a},\boldsymbol{\theta}) - I_{p,q}(\boldsymbol{\theta};\mathbf{a})\}$$

Intuition

The true joint distribution of (θ, a) is $p(\theta)\alpha(a)$

Fractions of population enjoy deviations to empirical distributions $q(\theta, a)$

For aggregate risk, we evaluate growth along typical state sequences

- ✗ no deviations of the marginal state distribution $q(\theta)$
- ✓ deviations of the correlation between a and θ

  fractions of population get "informed" by luck

Equivalence of the tradeoffs between

- growth advantage vs prevalence of deviation
- benefit vs cost of information

Sampled Choice Rule

Sample a $\$$ uniformly in a period t s.t. $\boldsymbol{\theta}_t = \theta$. This $\$$ originated from action a with probability

$$\tilde{q}(a \mid \theta) = \frac{\alpha^*(a) e^{u(a,\theta)}}{\mathsf{E}_{\alpha^*}\, e^{u(\mathbf{a},\theta)}}$$

Corollary

Sampled choice rule $=$ RI-optimal choice rule.

Statistician who samples proportionally to wealth, learns $\tilde{q}$

Successful individuals appear to be optimally informed

Wealth Concentration

Retrospective growth rate is $\mathsf{E}_{p,\boldsymbol{q}^*}\, u(\mathbf{a}, \boldsymbol{\theta})$

Growth rate of aggregate wealth is only $\mathsf{E}_{p,q^*}\, u(\mathbf{a}, \boldsymbol{\theta}) - I_{p,q^*}(\boldsymbol{\theta}; \mathbf{a})$

Fraction $e^{-t \times I_{p,q^*}(\boldsymbol{\theta};\mathbf{a})}$ of population owns “all” wealth in period t

information expense = rate of wealth concentration

Transfer of Methods

Growth-optimal portfolio problem is hard

- rational-inattention machinery may help

Proof of concept:

- two questions that require analysis of posteriors
- this is natural in the RI context
- but hard to think of in the growth context

Free Public Information

Aggregate states θ_t

Public signals y_t

Known joint distribution $p(\theta, y)$, iid across t

Signal-contingent hedging $\alpha(a \mid y)$

$$\max_{\alpha} \mathsf{E}_p \ln \left(\mathsf{E}_{\alpha(a|\mathbf{y})} \, e^{u(\mathbf{a},\boldsymbol{\theta})} \right)$$

$\alpha^*(a)$ and $\alpha^{**}(a \mid y)$ – optimal strategies without and with information

V^* and V^{**} – growth rates achieved without and with information

Impact of Side Information

$q^*(\theta \mid a)$ – optimal posterior without information

regularity condition

Weak signal: $p(\theta \mid y)$ are in the convex hull of $q^*(\theta \mid a)$, $a : q^*(a) > 0$.

Proposition

Provision of public signal increases growth rate by

$$V^{**} - V^* = I(\boldsymbol{\theta}; \mathbf{y}).$$

Marginal action distributions are unaffected:

$$\alpha^*(a) = \alpha^{**}(a).$$

Rate of wealth concentration decreases.

Proof

Posterior approach of Caplin&Dean&Leahy'22

$$\max_{\mathbf{r}} \quad \mathrm{E}\left[v(\mathbf{r}) + H(\mathbf{r}) - H(p)\right]$$
$$\text{s.t.:} \quad \mathrm{E}\,\mathbf{r} = p$$

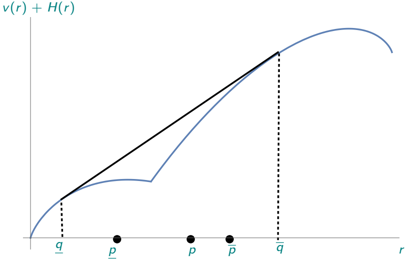

Proof

Posterior approach of Caplin&Dean&Leahy'22

$$\begin{array}{ll} \max_{\mathbf{r}} & \mathrm{E}\left[v(\mathbf{r}) + H(\mathbf{r}) - H(p)\right] \\ \text{s.t.:} & \mathrm{E}\,\mathbf{r} = p \end{array}$$

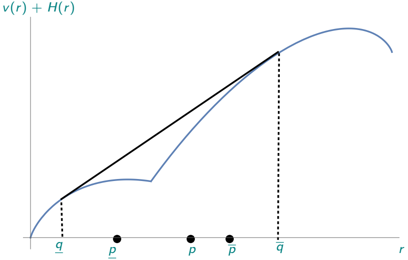

Martingale argument: free signal doesn't affect $\mathrm{E}\left[v(\mathbf{r}) + H(\mathbf{r})\right]$

Benefit of the free signal $=$ reduction of the prior entropy

Cost of Misspecification

Loss from misperceiving p as p'

$$L(p, p') = \mathrm{E}_p \ln \mathrm{E}_{\alpha^*(p)} e^{u(\mathbf{a},\boldsymbol{\theta})} - \mathrm{E}_p \ln \mathrm{E}_{\alpha^*(p')} e^{u(\mathbf{a},\boldsymbol{\theta})}$$

regularity condition

Small misperception: p' is in the convex hull of optimal posteriors.

Universal Cost of Misperception

$$L(p, p') = D\left(p \parallel p'\right).$$

Proof

Consider the value function $V^*(p) = \max_\alpha \mathrm{E}_p \ln \mathrm{E}_\alpha \, e^{u(\mathbf{a},\boldsymbol{\theta})}$

$V^*(p) = \underbrace{\mathrm{E}\left[v(\mathbf{r}^*) + H(\mathbf{r}^*)\right]}_{\text{linear in } p} - H(p)$ on the convex hull

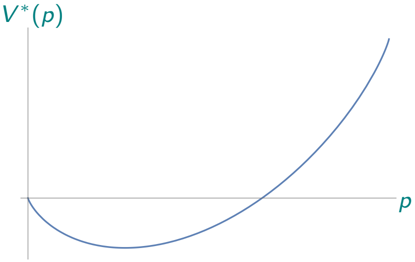

Proof

Graph of $\mathsf{E}_p \ln \mathsf{E}_{\alpha^*(p')} e^{u(\mathbf{a},\boldsymbol{\theta})}$ is a hyperplane tangent to $V^*(p)$ at $p = p'$

- linearity in p – it's an expectation
- tangency – optimality of $\alpha^*(p')$

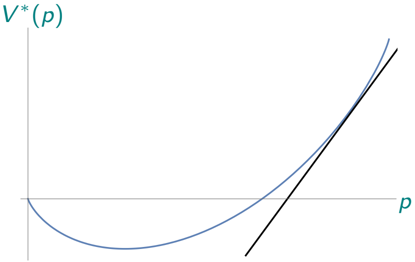

Proof

$L(p, p')$ is the error of linear approximation of V^*

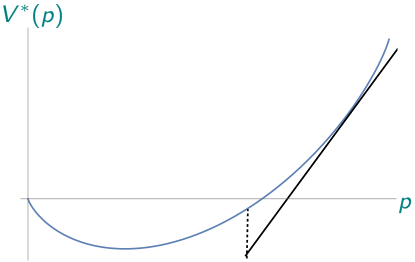

Proof

$L(p, p')$ is the error of linear approximation of $-H$

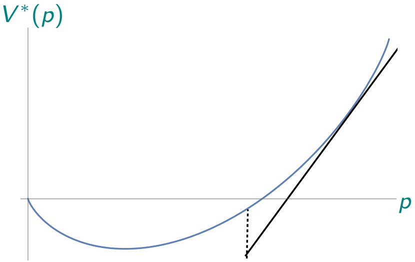

Proof

but that's $D(p \parallel p')$ (KL is a Bregman divergence)

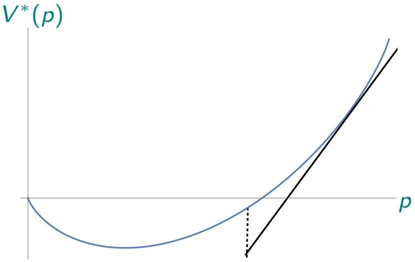

Period Length

Period length δ changes $u(a, \theta)$ to $\delta u(a, \theta)$

Idiosyncratic risk becomes

$$\max_{a,q} \left\{ \mathrm{E}_q\, u(a, \boldsymbol{\theta}) - \frac{1}{\delta} D(q \parallel p) \right\}$$

Aggregate risk becomes

$$\max_{q \in \Delta(A)^{\Theta}} \left\{ \mathrm{E}_{p,q}\, u(\mathbf{a}, \boldsymbol{\theta}) - \frac{1}{\delta} I_{p,q}(\mathbf{a}; \boldsymbol{\theta}) \right\}$$

Both cases collapse to EUT for (p, u) as $\delta \to 0$

Conclusion

Stochastic growth unfolds along a large deviation

Lots of inequality

Empirical implications: winners' data need careful interpretation

decision theory $\rightleftarrows$ growth theory

Transfers of insights across domains