REFERENCES

- Hansen, Beliefs, Doubts and Learning: Valuing Macroeconomic Risk, Ely, AER
- Hansen, Sargent, Tallarini, Robust Permanent Income and Pricing, ReStud
- Anderson, Hansen, Sargent, A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection, Journal of European Economic Association, 2003
- Maenhout, Robust Portfolio Rules and Asset Pricing, Review of Financial Studies, 2004.
- Maenhout, Robust Portfolio Rules and Detection Error Probabilities for a Mean Reverting Risk Premium, JET, 2006.
- Hansen, Sargent, Fragile Beliefs

Rational Expectations from Two Perspectives

- Econometricians - impose rational expectations - use cross equation restrictions that assume agent knowledge of parameters
- Economic decision makers - make investment decisions - forecast the future- arguably not knowing parameters

Should economic agents and econometricians be placed on a more equal footing, or not?

Rational Expectations from Two Perspectives

- Econometricians - impose rational expectations - use cross equation restrictions that assume agent knowledge of parameters
- Economic decision makers - make investment decisions - forecast the future- arguably not knowing parameters

Should economic agents and econometricians be placed on a more equal footing, or not?

Rational Expectations from Two Perspectives

- Econometricians - impose rational expectations - use cross equation restrictions that assume agent knowledge of parameters
- Economic decision makers - make investment decisions - forecast the future- arguably not knowing parameters

Should economic agents and econometricians be placed on a more equal footing, or not?

Rational Expectations from Two Perspectives

- Econometricians - impose rational expectations - use cross equation restrictions that assume agent knowledge of parameters
- Economic decision makers - make investment decisions - forecast the future- arguably not knowing parameters

Should economic agents and econometricians be placed on a more equal footing, or not?

Fundamental Questions

- When is estimation difficult?
- What are the consequences for the econometrician?
- What are the consequences for economic agents and for equilibrium outcomes?
- What are the real time consequences of learning for competitive security markets?
- How is learning altered when decision-makers admit that the models are misspecified or simplified?

Fundamental Questions

- When is estimation difficult?
- What are the consequences for the econometrician?
- What are the consequences for economic agents and for equilibrium outcomes?
- What are the real time consequences of learning for competitive security markets?
- How is learning altered when decision-makers admit that the models are misspecified or simplified?

Fundamental Questions

- When is estimation difficult?
- What are the consequences for the econometrician?
- What are the consequences for economic agents and for equilibrium outcomes?
- What are the real time consequences of learning for competitive security markets?
- How is learning altered when decision-makers admit that the models are misspecified or simplified?

Fundamental Questions

- When is estimation difficult?
- What are the consequences for the econometrician?
- What are the consequences for economic agents and for equilibrium outcomes?
- What are the real time consequences of learning for competitive security markets?
- How is learning altered when decision-makers admit that the models are misspecified or simplified?

Fundamental Questions

- When is estimation difficult?
- What are the consequences for the econometrician?
- What are the consequences for economic agents and for equilibrium outcomes?
- What are the real time consequences of learning for competitive security markets?
- How is learning altered when decision-makers admit that the models are misspecified or simplified?

Fundamental Questions

- When is estimation difficult?
- What are the consequences for the econometrician?
- What are the consequences for economic agents and for equilibrium outcomes?
- What are the real time consequences of learning for competitive security markets?
- How is learning altered when decision-makers admit that the models are misspecified or simplified?

Application: Macroeconomics and Asset Pricing

1. How does statistical ambiguity alter the predicted risk-return relation?

   Explains part of the steep slope by ...

2. Can learning induce model uncertainty premia that are larger when macroeconomic growth is sluggish?

   Yes, but only if ...

Application: Macroeconomics and Asset Pricing

1. How does statistical ambiguity alter the predicted risk-return relation?

   Explains part of the steep slope by ...

2. Can learning induce model uncertainty premia that are larger when macroeconomic growth is sluggish?

   Yes, but only if ...

Application: Macroeconomics and Asset Pricing

1. How does statistical ambiguity alter the predicted risk-return relation?

   Explains part of the steep slope by ...

2. Can learning induce model uncertainty premia that are larger when macroeconomic growth is sluggish?

   Yes, but only if ...

Rational Expectations Econometrics

Key words: Cross-equation restrictions
Key assumption: Investor knowledge

Illustrations from the asset pricing literature: risk prices.

Model ingredients:

- Consumption dynamics:

$$
\begin{aligned}
c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\
z_{t+1} &= A z_t + \sigma_z u_{t+1},
\end{aligned}
$$

  where $\{u_t\}$ is an iid sequence of normally distributed random vectors.

- - Kreps-Porteus, Epstein-Zin and others preferences in which the intertemporal composition of risk matters - Bansal-Yaron feature predictability.

Rational Expectations Econometrics

Key words: Cross-equation restrictions
Key assumption: Investor knowledge

Illustrations from the asset pricing literature: risk prices.

Model ingredients:

- Consumption dynamics:

$$
\begin{aligned}
c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\
z_{t+1} &= A z_t + \sigma_z u_{t+1},
\end{aligned}
$$

  where $\{u_t\}$ is an iid sequence of normally distributed random vectors.

- \- Kreps-Porteus, Epstein-Zin and others preferences in which the intertemporal composition of risk matters - Bansal-Yaron feature predictability.

RATIONAL EXPECTATIONS ECONOMETRICS

Key words: Cross-equation restrictions
Key assumption: Investor knowledge

Illustrations from the asset pricing literature: risk prices.

Model ingredients:

- Consumption dynamics:

$$
\begin{aligned}
c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\
z_{t+1} &= A z_t + \sigma_z u_{t+1},
\end{aligned}
$$

  where $\{u_t\}$ is an iid sequence of normally distributed random vectors.

- - Kreps-Porteus, Epstein-Zin and others preferences in which the intertemporal composition of risk matters - Bansal-Yaron feature predictability.

Rational Expectations Econometrics

Key words: Cross-equation restrictions
Key assumption: Investor knowledge

Illustrations from the asset pricing literature: risk prices.

Model ingredients:

- Consumption dynamics:

$$\begin{aligned} c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\ z_{t+1} &= A z_t + \sigma_z u_{t+1}, \end{aligned}$$

  where $\{u_t\}$ is an iid sequence of normally distributed random vectors.
- \- Kreps-Porteus, Epstein-Zin and others preferences in which the intertemporal composition of risk matters - Bansal-Yaron feature predictability.

Risk prices

Recall

$$\begin{aligned} c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\ z_{t+1} &= A z_t + \sigma_z u_{t+1}, \end{aligned}$$

Assume a IES = 1, a recursive utility risk parameter γ and a discount factor β.

Price the one-period exposure to shock u_{t+1} with a known distribution. Prices are quoted in terms of mean reward.

$$p = \sigma_c + (\gamma - 1)\left[\beta\alpha(I - \beta A)^{-1}\sigma_z + \sigma_c\right]$$

Limiting long horizon risk prices:

$$p_\infty = \left[\sigma_c + \alpha(I - A)^{-1}\sigma_z\right] + (\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

Cross equation restrictions link the consumption dynamics and the risk prices.

RISK PRICES

Recall

$$\begin{aligned} c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\ z_{t+1} &= A z_t + \sigma_z u_{t+1}, \end{aligned}$$

Assume a IES = 1, a recursive utility risk parameter γ and a discount factor β.

Price the one-period exposure to shock u_{t+1} with a known distribution. Prices are quoted in terms of mean reward.

$$p = \sigma_c + (\gamma - 1)\left[\beta\alpha(I - \beta A)^{-1}\sigma_z + \sigma_c\right]$$

Limiting long horizon risk prices:

$$p_\infty = \left[\sigma_c + \alpha(I - A)^{-1}\sigma_z\right] + (\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

Cross equation restrictions link the consumption dynamics and the risk prices.

Risk prices

Recall

$$\begin{aligned} c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\ z_{t+1} &= A z_t + \sigma_z u_{t+1}, \end{aligned}$$

Assume a IES = 1, a recursive utility risk parameter γ and a discount factor β.

Price the one-period exposure to shock u_{t+1} with a known distribution. Prices are quoted in terms of mean reward.

$$p = \sigma_c + (\gamma - 1) \left[\beta\alpha(I - \beta A)^{-1}\sigma_z + \sigma_c \right]$$

Limiting long horizon risk prices:

$$p_\infty = \left[\sigma_c + \alpha(I - A)^{-1}\sigma_z \right] + (\gamma - 1) \left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z \right]$$

Cross equation restrictions link the consumption dynamics and the risk prices.

Econometrics and Limited Information

Asset pricing implications represented as conditional moment restrictions conditioned on investors information.

Apply Law of Iterated Expectations to deduce corresponding unconditional moment restrictions.

Hansen-Singleton, Hansen-Richard, Luttmer and others.

Exploit the potential information advantage of investors in deducing testable restrictions.

Statistical ambiguity

Question: How does statistical ambiguity alter the predicted risk-return relation?

Two perspectives:

- Econometrician
- Economic agents

When is Statistical Inference Challenging

Problem: Suppose there are two models under consideration: model a and model b. Historical data are available to select the correct model.

Chernoff: Pose a simple decision problem and ask how likely is it to make a mistake? What is the decay rate of the mistake probabilities per unit of observation?

Details

- average errors or equate errors
- compare log-likelihood ratio to a threshold

When is Statistical Inference Challenging

Problem: Suppose there are two models under consideration: model a and model b. Historical data are available to select the correct model.

Chernoff: Pose a simple decision problem and ask how likely is it to make a mistake? What is the decay rate of the mistake probabilities per unit of observation?

Details

- average errors or equate errors
- compare log-likelihood ratio to a threshold

When is Statistical Inference Challenging

Problem: Suppose there are two models under consideration: model a and model b. Historical data are available to select the correct model.

Chernoff: Pose a simple decision problem and ask how likely is it to make a mistake? What is the decay rate of the mistake probabilities per unit of observation?

Details

- average errors or equate errors
- compare log-likelihood ratio to a threshold

Example with normal data

Two models

A) mean μ_a and variance Σ.

B) mean μ_b with variance Σ.

Mistake probabilities eventually decay as function of sample size at rate:

$$\frac{1}{8}(\mu_a - \mu_b)'\Sigma^{-1}(\mu_a - \mu_b)$$

Call this the Chernoff rate.

EXAMPLE WITH NORMAL DATA

Two models

A) mean μ_a and variance Σ.

B) mean μ_b with variance Σ.

Mistake probabilities eventually decay as function of sample size at rate:

$$\frac{1}{8}(\mu_a - \mu_b)'\Sigma^{-1}(\mu_a - \mu_b)$$

Call this the Chernoff rate.

EXAMPLE WITH NORMAL DATA

Two models

A) mean μ_a and variance Σ.

B) mean μ_b with variance Σ.

Mistake probabilities eventually decay as function of sample size at rate:

$$\frac{1}{8}(\mu_a - \mu_b)'\Sigma^{-1}(\mu_a - \mu_b)$$

Call this the Chernoff rate.

How Close are Models Statistically?

We can ask this question in a more general context.

- Does not require normality; study behavior of relative likelihoods.
- Does not require independence; Can be extended to Markov processes.
- Can average mistake probabilities or equate them.
- The (overly) simplified problem of making pairwise comparison between models remains informative, as we will see.

How Close are Models Statistically?

We can ask this question in a more general context.

- Does not require normality; study behavior of relative likelihoods.
- Does not require independence; Can be extended to Markov processes.
- Can average mistake probabilities or equate them.
- The (overly) simplified problem of making pairwise comparison between models remains informative, as we will see.

How Close are Models Statistically?

We can ask this question in a more general context.

- Does not require normality; study behavior of relative likelihoods.
- Does not require independence; Can be extended to Markov processes.
- Can average mistake probabilities or equate them.
- The (overly) simplified problem of making pairwise comparison between models remains informative, as we will see.

How Close are Models Statistically?

We can ask this question in a more general context.

- Does not require normality; study behavior of relative likelihoods.
- Does not require independence; Can be extended to Markov processes.
- Can average mistake probabilities or equate them.
- The (overly) simplified problem of making pairwise comparison between models remains informative, as we will see.

How Close are Models Statistically?

We can ask this question in a more general context.

- Does not require normality; study behavior of relative likelihoods.
- Does not require independence; Can be extended to Markov processes.
- Can average mistake probabilities or equate them.
- The (overly) simplified problem of making pairwise comparison between models remains informative, as we will see.

A Challenge to an Econometrician

Model a has the form:

$$\begin{aligned} c_{t+1} - c_t &= .0058 + z_t + .0053 u_{1,t+1} \\ z_{t+1} &= .98 z_t + .00025 u_{2,t+1}. \end{aligned}$$

Another representation of z_t:

$$z_t = .00025 \sum_{j=0}^{\infty} (.98)^j u_{2,t-j}.$$

Illustrates a model of Bansal-Yaron: low frequency component to consumption predictability.

Model b has the same form but the second shock is eliminated.

Mistake Probabilities for Consumption Dynamics

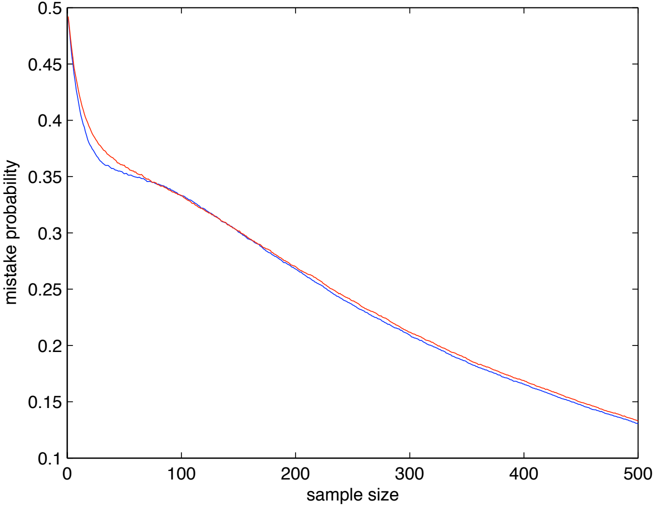

Mistake probability as a function of sample size for the predictable

Logarithm of Mistake Probabilities

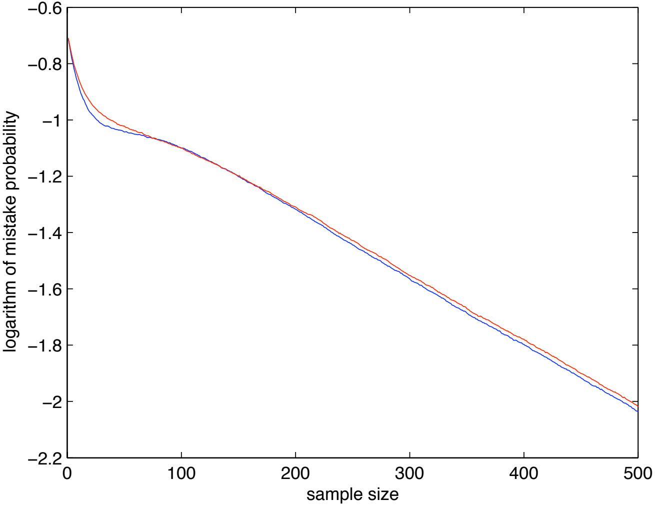

Logarithm of mistake probability as a function of sample size for the predictable growth model vis a vis the iid growth model.

PRIORS AND POSTERIORS

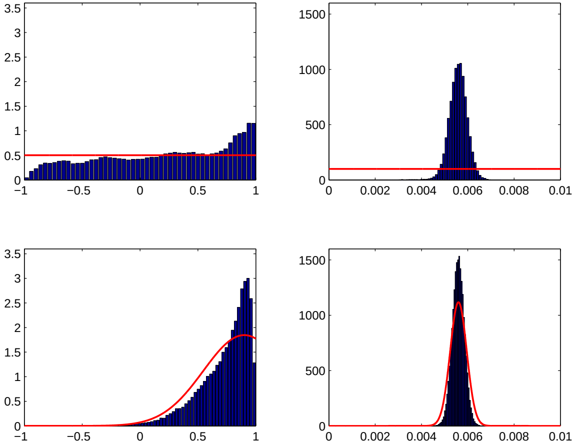

Left panel is the AR parameter for the hidden state; right panel is the mean growth rate of consumption; red line is the prior.

Review of risk prices

- Represent risks as $\mu + \Lambda u$ where u is a random vector with mean zero and an identity as its covariance matrix.
- The covariance matrix of the implied risks is $\Sigma = \Lambda\Lambda'$. ’
- A asset pricing model restricts the mean return vector μ as a function of the risk exposure Λ by assigning a risk price vector p:

  $$\mu - r^f \mathbf{1}_n = \Lambda p$$

  where r^f is the return on a risk free asset and $\mathbf{1}_n$ is an n dimensional vector of ones.

Review of risk prices

- Represent risks as $\mu + \Lambda u$ where u is a random vector with mean zero and an identity as its covariance matrix.
- The covariance matrix of the implied risks is $\Sigma = \Lambda\Lambda'$. '
- A asset pricing model restricts the mean return vector μ as a function of the risk exposure Λ by assigning a risk price vector p:

$$\mu - r^f \mathbf{1}_n = \Lambda p$$

where r^f is the return on a risk free asset and $\mathbf{1}_n$ is an n dimensional vector of ones.

Risk-return tradeoff

To summarize asset return implications, maximize Sharpe ratios of portfolios by choice of the weight vector ω:

$$
\begin{aligned}
\max_{\omega} \frac{\omega \cdot (\mu - \mathbf{1}_n r^f)}{\sqrt{\omega' \Sigma \omega}} &= \max_{\omega} \frac{\omega' \Lambda p}{\sqrt{\omega' \Sigma \omega}} \\
&= |p| \\
&= \left[(\mu - \mathbf{1}_n r^f)' \Sigma^{-1} (\mu - \mathbf{1}_n r^f) \right]^{1/2}.
\end{aligned}
$$

RISK PRICES AND PUZZLES

Observations

- Risk prices - assign prices to the exposure to alternative shocks
- Risk return tradeoff - Sharpe ratios - implied by risk prices

Risk prices are of direct interest as a challenge for an asset pricing model when they can be measured. A weaker challenge is compare to lower bounds on the risk return tradeoff.

A steep risk-return tradeoff is a challenge for asset pricing models without appealing to high risk aversion.

RISK PRICES AND PUZZLES

Observations

- Risk prices - assign prices to the exposure to alternative shocks
- Risk return tradeoff - Sharpe ratios - implied by risk prices

Risk prices are of direct interest as a challenge for an asset pricing model when they can be measured. A weaker challenge is compare to lower bounds on the risk return tradeoff.

A steep risk-return tradeoff is a challenge for asset pricing models without appealing to high risk aversion.

Can a small amount of statistical ambiguity explain a steep risk-return frontier?

Recall, maximum Sharpe ratio:

$$|p| = \left[(\mu - \mathbf{1}_n r^f)' \Sigma^{-1} (\mu - \mathbf{1}_n r^f)\right]^{1/2} \tag{1}$$

When $|p|$ based on the hypothetically correct value μ is small, changing μ to $\tilde{\mu}$ in (1) is dominated by:

$$\left[(\mu - \tilde{\mu})' \Sigma^{-1} (\mu - \tilde{\mu})\right]^{1/2} \tag{2}$$

Squaring (2) and dividing by eight gives the Chernoff rate:

$$\frac{(\mu - \tilde{\mu})' \Sigma^{-1} (\mu - \tilde{\mu})}{8}$$

which quantifies statistical latitude in moving μ!

Can a small amount of statistical ambiguity explain a steep risk-return frontier?

Recall, maximum Sharpe ratio:

$$|p| = \left[(\mu - \mathbf{1}_n r^f)' \Sigma^{-1} (\mu - \mathbf{1}_n r^f)\right]^{1/2} \tag{1}$$

When $|p|$ based on the hypothetically correct value μ is small, changing μ to $\tilde{\mu}$ in (1) is dominated by:

$$\left[(\mu - \tilde{\mu})' \Sigma^{-1} (\mu - \tilde{\mu})\right]^{1/2} \tag{2}$$

Squaring (2) and dividing by eight gives the Chernoff rate:

$$\frac{(\mu - \tilde{\mu})' \Sigma^{-1} (\mu - \tilde{\mu})}{8}$$

which quantifies statistical latitude in moving μ!

CHERNOFF RATE $\Longleftrightarrow$ PRICE OF UNCERTAINTY

Let us quantify the link between the statistical discrimination of alternative models and the risk prices as depicted in empirical finance.

As a rough idea, consider a Chernoff rate per observation of .13$\%$ quarterly or about .5$\%$ on an annual basis. This changes the quarterly Sharpe ratio by about .1 for quarterly data.

How do we interpret this movement?

- Statistical uncertainty from the standpoint of the econometrician ;
- Aversion or statistical ambiguity on the part of investors;

Extend these ideas to dynamic, nonlinear Markov environments in a formal way. (Anderson-Hansen-Sargent.)

CHERNOFF RATE $\Longleftrightarrow$ PRICE OF UNCERTAINTY

Let us quantify the link between the statistical discrimination of alternative models and the risk prices as depicted in empirical finance.

As a rough idea, consider a Chernoff rate per observation of $.13\%$ quarterly or about $.5\%$ on an annual basis. This changes the quarterly Sharpe ratio by about .1 for quarterly data.

How do we interpret this movement?

- Statistical uncertainty from the standpoint of the econometrician ;
- Aversion or statistical ambiguity on the part of investors;

Extend these ideas to dynamic, nonlinear Markov environments in a formal way. (Anderson-Hansen-Sargent.)

CHERNOFF RATE $\Longleftrightarrow$ PRICE OF UNCERTAINTY

Let us quantify the link between the statistical discrimination of alternative models and the risk prices as depicted in empirical finance.

As a rough idea, consider a Chernoff rate per observation of .13% quarterly or about .5% on an annual basis. This changes the quarterly Sharpe ratio by about .1 for quarterly data.

How do we interpret this movement?

- Statistical uncertainty from the standpoint of the econometrician ;
- Aversion or statistical ambiguity on the part of investors;

Extend these ideas to dynamic, nonlinear Markov environments in a formal way. (Anderson-Hansen-Sargent.)

CHERNOFF RATE $\Longleftrightarrow$ PRICE OF UNCERTAINTY

Let us quantify the link between the statistical discrimination of alternative models and the risk prices as depicted in empirical finance.

As a rough idea, consider a Chernoff rate per observation of .13% quarterly or about .5% on an annual basis. This changes the quarterly Sharpe ratio by about .1 for quarterly data.

How do we interpret this movement?

- Statistical uncertainty from the standpoint of the econometrician ;
- Aversion or statistical ambiguity on the part of investors;

Extend these ideas to dynamic, nonlinear Markov environments in a formal way. (Anderson-Hansen-Sargent.)

Consumption growth model reconsidered

- Consumption dynamics:

$$\begin{aligned} c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\ z_{t+1} &= A z_t + \sigma_z u_{t+1}, \end{aligned}$$

- One period risk prices

$$p = \sigma_c + (\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

- Statistical analysis of

$$(\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

suggests an alternative interpretation. Reinterpret $\gamma - 1$ as reflecting statistical ambiguity instead of risk aversion. For instance, γ of about 8.5 gives the .5% decay rate per annum.

Consumption growth model reconsidered

- Consumption dynamics:

$$\begin{aligned} c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\ z_{t+1} &= A z_t + \sigma_z u_{t+1}, \end{aligned}$$

- One period risk prices

$$p = \sigma_c + (\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

- Statistical analysis of

$$(\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

suggests an alternative interpretation. Reinterpret $\gamma - 1$ as reflecting statistical ambiguity instead of risk aversion. For instance, γ of about 8.5 gives the .5% decay rate per annum.

Consumption growth model reconsidered

- Consumption dynamics:

$$\begin{aligned} c_{t+1} - c_t &= \mu_c + \alpha z_t + \sigma_c u_{t+1} \\ z_{t+1} &= A z_t + \sigma_z u_{t+1}, \end{aligned}$$

- One period risk prices

$$p = \sigma_c + (\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

- Statistical analysis of

$$(\gamma - 1)\left[\sigma_c + \beta\alpha(I - \beta A)^{-1}\sigma_z\right]$$

suggests an alternative interpretation. Reinterpret $\gamma - 1$ as reflecting statistical ambiguity instead of risk aversion. For instance, γ of about 8.5 gives the .5% decay rate per annum.

Transition to Learning

The Chernoff analysis characterizes our ability to tell models apart statistically. It studies how potent learning can be, but it has little to say about the dynamics of learning.

Next I want to consider the real time impact of learning when there are only weak signals about some aspects of the economic environment.

There is empirical evidence that risk premia are larger in macroeconomic downturns. (Campbell and Cochrane) and others.

Can learning induce model uncertainty premia that are larger when macroeconomic growth is more sluggish?

Transition to Learning

The Chernoff analysis characterizes our ability to tell models apart statistically. It studies how potent learning can be, but it has little to say about the dynamics of learning.

Next I want to consider the real time impact of learning when there are only weak signals about some aspects of the economic environment.

There is empirical evidence that risk premia are larger in macroeconomic downturns. (Campbell and Cochrane) and others.

Can learning induce model uncertainty premia that are larger when macroeconomic growth is more sluggish?

Transition to Learning

The Chernoff analysis characterizes our ability to tell models apart statistically. It studies how potent learning can be, but it has little to say about the dynamics of learning.

Next I want to consider the real time impact of learning when there are only weak signals about some aspects of the economic environment.

There is empirical evidence that risk premia are larger in macroeconomic downturns. (Campbell and Cochrane) and others.

Can learning induce model uncertainty premia that are larger when macroeconomic growth is more sluggish?

Transition to Learning

The Chernoff analysis characterizes our ability to tell models apart statistically. It studies how potent learning can be, but it has little to say about the dynamics of learning.

Next I want to consider the real time impact of learning when there are only weak signals about some aspects of the economic environment.

There is empirical evidence that risk premia are larger in macroeconomic downturns. (Campbell and Cochrane) and others.

Can learning induce model uncertainty premia that are larger when macroeconomic growth is more sluggish?

Learning

Bray-Kreps taxonomy distinguishes between

- Learning about a rational expectations equilibrium

  Uses learning rules with at least temporary misspecification, but without agents addressing this misspecification.
  There is no scope for uncertainty premia.
- Learning within a rational expectations equilibrium

I will explore learning within an equilibrium but include a concern for misspecification.

Learning

Bray-Kreps taxonomy distinguishes between

- Learning about a rational expectations equilibrium

  Uses learning rules with at least temporary misspecification, but without agents addressing this misspecification.
  There is no scope for uncertainty premia.
- Learning within a rational expectations equilibrium

I will explore learning within an equilibrium but include a concern for misspecification.

LEARNING AND THREE RELATED CONCEPTS

Outcome or signal s^* that depends on a function of a hidden state z. This hidden state summarizes all information potentially pertinent for characterizing the signal distribution.

Repeated over time as in Hidden State Markov Model (HMM).

Let $\mathcal{H}$ denote the history and current and past signals. Compute the distribution for z and hence s^* conditioned on $\mathcal{H}$.

- Law of iterated expectations
- Reduction of compound Lotteries
- Filtering - recursive implementation

Filtering

Filtering is a recursive way to reduce the lottery by averaging over the hidden state z.
Consider a signal:

$$ds_t = \kappa \cdot z_t dt + \sigma dB_t$$

where

- $\{z_t\}$ be a hidden state Markov chain. $\exp(tA)$ is the transition matrix over an interval of time t.
- Realized value of z_t is a coordinate vector. $\kappa \cdot z_t$ selects randomly among the entries in the vector κ.
- dB_t is a Brownian increment conditioned on the state z_t.

Special case of a Hamilton regime shift model. Used by David and Veronesi in asset pricing literature.

Filtering

Filtering is a recursive way to reduce the lottery by averaging over the hidden state z.
Consider a signal:

$$ds_t = \kappa \cdot z_t dt + \sigma dB_t$$

where

- $\{z_t\}$ be a hidden state Markov chain. $\exp(tA)$ is the transition matrix over an interval of time t.
- Realized value of z_t is a coordinate vector. $\kappa \cdot z_t$ selects randomly among the entries in the vector κ.
- dB_t is a Brownian increment conditioned on the state z_t.

Special case of a Hamilton regime shift model. Used by David and Veronesi in asset pricing literature.

Filtering

Filtering is a recursive way to reduce the lottery by averaging over the hidden state z.
Consider a signal:

$$ds_t = \kappa \cdot z_t dt + \sigma dB_t$$

where

- $\{z_t\}$ be a hidden state Markov chain. $\exp(tA)$ is the transition matrix over an interval of time t.
- Realized value of z_t is a coordinate vector. $\kappa \cdot z_t$ selects randomly among the entries in the vector κ.
- dB_t is a Brownian increment conditioned on the state z_t.

Special case of a Hamilton regime shift model. Used by David and Veronesi in asset pricing literature.

FILTERING SOLUTION

Wonham gives the solution based on *reducing* compound lotteries while *updating* probabilities based on past data.

Let $\bar{z}_t = E\left(z_t | \mathcal{H}_t\right)$, which is the vector of hidden state probabilities. Aim is to compute $\bar{z}_t$ recursively.

The recursive solution is a stochastic differential equation: represented in terms of an alternative standard Brownian motion $\{\bar{B}_t\}$:

$$\begin{aligned} ds_t &= \kappa \cdot \bar{z}_t dt + \sigma d\bar{B}_t \\ d\bar{z}_t &= A' \bar{z}_t dt + \frac{1}{\sigma^2} \left[\text{diag}\{\bar{z}_t\}\right] \left(\kappa - \mathbf{1}_n \kappa \cdot \bar{z}_t\right) \left(ds_t - \kappa \cdot \bar{z}_t\right). \end{aligned}$$

Observations

- Dynamics determined by the underlying intensity matrix A.
- Local volatility for the signal has the same magnitude as that prior to reduction.

Filtering solution

Wonham gives the solution based on *reducing* compound lotteries while *updating* probabilities based on past data.

Let $\bar{z}_t = E\left(z_t | \mathcal{H}_t\right)$, which is the vector of hidden state probabilities. Aim is to compute $\bar{z}_t$ recursively.

The recursive solution is a stochastic differential equation: represented in terms of an alternative standard Brownian motion $\{\bar{B}_t\}$:

$$
\begin{aligned}
ds_t &= \kappa \cdot \bar{z}_t dt + \sigma d\bar{B}_t \\
d\bar{z}_t &= A'\bar{z}_t dt + \frac{1}{\sigma^2}\left[\text{diag}\{\bar{z}_t\}\right]\left(\kappa - \mathbf{1}_n \kappa \cdot \bar{z}_t\right)\left(ds_t - \kappa \cdot \bar{z}_t\right).
\end{aligned}
$$

Observations

- Dynamics determined by the underlying intensity matrix A.
- Local volatility for the signal has the same magnitude as that prior to reduction.

Filtering solution

Wonham gives the solution based on *reducing* compound lotteries while *updating* probabilities based on past data.

Let $\bar{z}_t = E\left(z_t | \mathcal{H}_t\right)$, which is the vector of hidden state probabilities. Aim is to compute $\bar{z}_t$ recursively.

The recursive solution is a stochastic differential equation: represented in terms of an alternative standard Brownian motion $\{\bar{B}_t\}$:

$$\begin{aligned} ds_t &= \kappa \cdot \bar{z}_t dt + \sigma d\bar{B}_t \\ d\bar{z}_t &= A'\bar{z}_t dt + \frac{1}{\sigma^2}\left[\text{diag}\{\bar{z}_t\}\right]\left(\kappa - \mathbf{1}_n \kappa \cdot \bar{z}_t\right)\left(ds_t - \kappa \cdot \bar{z}_t\right). \end{aligned}$$

Observations

- Dynamics determined by the underlying intensity matrix A.
- Local volatility for the signal has the same magnitude as that prior to reduction.

Filtering solution

Wonham gives the solution based on *reducing* compound lotteries while *updating* probabilities based on past data.

Let $\bar{z}_t = E\left(z_t|\mathcal{H}_t\right)$, which is the vector of hidden state probabilities. Aim is to compute $\bar{z}_t$ recursively.

The recursive solution is a stochastic differential equation: represented in terms of an alternative standard Brownian motion $\{\bar{B}_t\}$:

$$\begin{aligned} ds_t &= \kappa \cdot \bar{z}_t dt + \sigma d\bar{B}_t \\ d\bar{z}_t &= A'\bar{z}_t dt + \frac{1}{\sigma^2}\left[\text{diag}\{\bar{z}_t\}\right]\left(\kappa - \mathbf{1}_n \kappa \cdot \bar{z}_t\right)\left(ds_t - \kappa \cdot \bar{z}_t\right). \end{aligned}$$

Observations

- Dynamics determined by the underlying intensity matrix A.
- Local volatility for the signal has the same magnitude as that prior to reduction.

Consumption Dynamics and Risk Prices

Let consumption growth be the signal.

$$\begin{aligned} dc_t &= \kappa \cdot z_t + \sigma dB_t \\ &= \kappa \cdot \bar{z}_t + [\sigma dB_t + \kappa \cdot (z_t - \bar{z}_t) dt] \end{aligned}$$

- State estimation error is hidden in the local evolution of consumption.
- Local risk prices are the same under usual expected utility model. Long run prices are altered.

Alternatives to Expected Utility Preferences

- Intertemporal composition of risk matters. Kreps-Porteus
- - Distinguish between risk conditioned on z and uncertainty about z. Segal and Klibanoff-Marinacci-Mukerji
- - Uncertainty aversion or robustness - Gilboa-Schmeidler, Epstein-Schneider, Hansen-Sargent and others

I use preferences that are represented by minimizing over families of probability models subject to penalization. Maccheroni, Marinacci and Rustichini.

Exponential tilting of probabilities. Jacobson and Whittle

Alternatives to Expected Utility Preferences

- Intertemporal composition of risk matters. Kreps-Porteus
- - Distinguish between risk conditioned on z and uncertainty about z. Segal and Klibanoff-Marinacci-Mukerji
- - Uncertainty aversion or robustness - Gilboa-Schmeidler, Epstein-Schneider, Hansen-Sargent and others

I use preferences that are represented by minimizing over families of probability models subject to penalization. Maccheroni, Marinacci and Rustichini.

Exponential tilting of probabilities. Jacobson and Whittle

Alternatives to Expected Utility Preferences

- Intertemporal composition of risk matters. Kreps-Porteus
- \- Distinguish between risk conditioned on z and uncertainty about z. Segal and Klibanoff-Marinacci-Mukerji
- \- Uncertainty aversion or robustness - Gilboa-Schmeidler, Epstein-Schneider, Hansen-Sargent and others

I use preferences that are represented by minimizing over families of probability models subject to penalization. Maccheroni, Marinacci and Rustichini.

Exponential tilting of probabilities. Jacobson and Whittle

Alternatives to Expected Utility Preferences

- Intertemporal composition of risk matters. Kreps-Porteus
- - Distinguish between risk conditioned on z and uncertainty about z. Segal and Klibanoff-Marinacci-Mukerji
- - Uncertainty aversion or robustness - Gilboa-Schmeidler, Epstein-Schneider, Hansen-Sargent and others

I use preferences that are represented by minimizing over families of probability models subject to penalization. Maccheroni, Marinacci and Rustichini.

Exponential tilting of probabilities. Jacobson and Whittle

Alternatives to Expected Utility Preferences

- Intertemporal composition of risk matters. Kreps-Porteus
- - Distinguish between risk conditioned on z and uncertainty about z. Segal and Klibanoff-Marinacci-Mukerji
- - Uncertainty aversion or robustness - Gilboa-Schmeidler, Epstein-Schneider, Hansen-Sargent and others

I use preferences that are represented by minimizing over families of probability models subject to penalization. Maccheroni, Marinacci and Rustichini.

Exponential tilting of probabilities. Jacobson and Whittle

Equilibrium calculation

Imitate rational expectations approach.

- fictitious social planner - compute value functions and exponentially slanted probabilities based on these functions.
- The probability distortions are the uncertainty premia in a decentralized model with security markets.

Possible interpretation of probability distortions - computational device for risk premia - alternative beliefs - statistical ambiguity

Equilibrium calculation

Imitate rational expectations approach.

- fictitious social planner - compute value functions and exponentially slanted probabilities based on these functions.
- The probability distortions are the uncertainty premia in a decentralized model with security markets.

Possible interpretation of probability distortions - computational device for risk premia - alternative beliefs - statistical ambiguity

Robustness and Exponential Tilting

Planner with continuation values conditioned on a hidden state, say $v_i + c_t$ for $i = 1, 2, ..., n$.

Tilt probabilities towards the states with the smallest continuation values

$$v_i^* = \exp\left(-\frac{v_i}{\theta}\right)$$

for some positive value of the parameter θ. Large values of θ make v_i^* close to their constant value of unity.

The exponentially tilted probabilities are:

$$\tilde{z}_t^i = \frac{v_i^* \bar{z}_{i,t}}{\sum_i v_i^* \bar{z}_{i,t}}.$$

Weight more heavily the low continuation values.

Robustness and Exponential Tilting

Planner with continuation values conditioned on a hidden state, say $v_i + c_t$ for $i = 1, 2, ..., n$.

Tilt probabilities towards the states with the smallest continuation values

$$v_i^* = \exp\left(-\frac{v_i}{\theta}\right)$$

for some positive value of the parameter θ. Large values of θ make v_i^* close to their constant value of unity.

The exponentially tilted probabilities are:

$$\tilde{z}_t^i = \frac{v_i^* \bar{z}_{i,t}}{\sum_i v_i^* \bar{z}_{i,t}}.$$

Weight more heavily the low continuation values.

Robustness and Exponential Tilting

Planner with continuation values conditioned on a hidden state, say $v_i + c_t$ for $i = 1, 2, ..., n$.

Tilt probabilities towards the states with the smallest continuation values

$$v_i^* = \exp\left(-\frac{v_i}{\theta}\right)$$

for some positive value of the parameter θ. Large values of θ make v_i^* close to their constant value of unity.

The exponentially tilted probabilities are:

$$\tilde{z}_t^i = \frac{v_i^* \bar{z}_{i,t}}{\sum_i v_i^* \bar{z}_{i,t}}.$$

Weight more heavily the low continuation values.

Two questions

1. Where do the continuation values come from?

   Answer Solve Bellman equation. (Intertemporal elasticity is unity for simplicity.) Include additional risk adjustment or robustness adjustment

2. Where does the exponential tilting come from:

   Two alternative answers

   A) A smooth ambiguity adjustment made to the continuation value.
   B) Explore mistakes in the outcome of the filtering solution.

Two questions

1. Where do the continuation values come from?

   Answer Solve Bellman equation. (Intertemporal elasticity is unity for simplicity.) Include additional risk adjustment or robustness adjustment

2. Where does the exponential tilting come from:

   Two alternative answers

   A) A smooth ambiguity adjustment made to the continuation value.
   B) Explore mistakes in the outcome of the filtering solution.

Two questions

1. Where do the continuation values come from?

   Answer Solve Bellman equation. (Intertemporal elasticity is unity for simplicity.) Include additional risk adjustment or robustness adjustment

2. Where does the exponential tilting come from:

   Two alternative answers

   A) A smooth ambiguity adjustment made to the continuation value.
   B) Explore mistakes in the outcome of the filtering solution.

Two questions

1. Where do the continuation values come from?

   Answer Solve Bellman equation. (Intertemporal elasticity is unity for simplicity.) Include additional risk adjustment or robustness adjustment

2. Where does the exponential tilting come from:

   Two alternative answers

   A) A smooth ambiguity adjustment made to the continuation value.
   B) Explore mistakes in the outcome of the filtering solution.

Two questions

1. Where do the continuation values come from?

   Answer Solve Bellman equation. (Intertemporal elasticity is unity for simplicity.) Include additional risk adjustment or robustness adjustment

2. Where does the exponential tilting come from:

   Two alternative answers

   A) A smooth ambiguity adjustment made to the continuation value.
   B) Explore mistakes in the outcome of the filtering solution.

Two alternative formulations with Exponential Tilting

1. Concern about misspecified dynamics for consumption and for the probability updating. Recursive utility.
2. Separate concerns about misspecified dynamics and misspecified state probabilities. Include a sensitivity analysis to priors.

Both include versions of exponential tilting and both use the filtering solution as a benchmark. Second has two separate adjustments.

Two alternative formulations with Exponential Tilting

1. Concern about misspecified dynamics for consumption and for the probability updating. Recursive utility.
2. Separate concerns about misspecified dynamics and misspecified state probabilities. Include a sensitivity analysis to priors.

Both include versions of exponential tilting and both use the filtering solution as a benchmark. Second has two separate adjustments.

Risk and Uncertainty Premia I

$$\sigma + \frac{\sigma}{\theta_f} + \frac{1}{\theta_f}\Delta(\bar{z}) \cdot \frac{\partial V(\bar{z})}{\partial \bar{z}}$$

Ex Utility IES $= 1$	exponential tilting full information	exponential tilting state estimation
σ	$\frac{\sigma}{\theta_f}$	$\frac{1}{\theta_f}\Delta(\bar{z}) \cdot \frac{\partial V(\bar{z})}{\partial \bar{z}}$
time invariant	time invariant	time varying

value function: $V(\bar{z}) + c$;

σ: response of consumption to new information;

$\Delta(\bar{z})$: vector response of the probabilities to new information.

Uncertainty Premia as Function of Probability

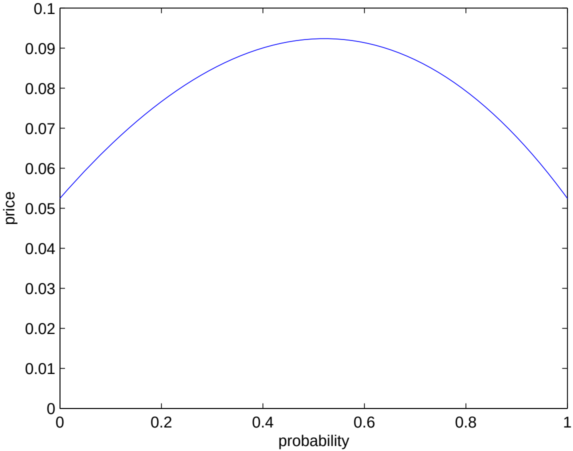

Peak impact at the point in which filtered probability is one half.
Cagetti, Hansen, Sargent and Williams.

Risk and Uncertainty Prices

$$\sigma + \frac{\sigma}{\theta_f} + \left[\frac{1}{\sigma} \left(\bar{z}_t - \tilde{z}_t \right) \cdot \kappa \right]$$

Hansen-Sargent - Fragile Beliefs and the Price of Model Uncertainty

Ex Utility IES is one	exponential tilting state dynamics	exponential tilting state probabilities
σ time invariant	$\frac{\sigma}{\theta_f}$ time invariant	$\frac{(\bar{z}-\tilde{z})\cdot\kappa}{\sigma}$ time varying

value function: $V(z) + c$;

σ: response of consumption to new information,

$\tilde{z}$: is the exponentially tilted counterpart of $\bar{z}$ probabilities;

κ: vector of alternative growth rates.

Learning Dynamics of Uncertainty Premia

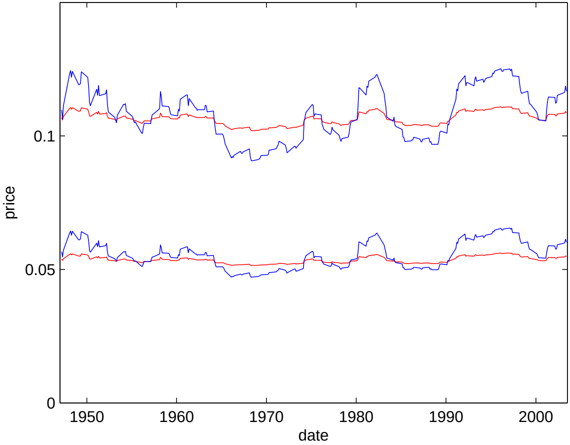

Agents are selecting among two models and learning about parameters. Hansen and Sargent

Source of Variation in Uncertainty Premia

- Induced by probability slanting and hence relative magnitudes of the continuation values;
- Asymmetry
  - i) the average continuation value for the model with predictable consumption growth is lower than the model without predictability - extra channel for model misspecification
  - ii) good consumption growth realizations, expected consumption growth is higher for the model with predictable consumption growth and the relative magnitudes are closer together.
  - iii) bad consumption growth realizations, expected consumption growth lower for model with predictable consumption growth and relative magnitudes across models are further away.