A retrieved-context theory of financial decisions

Jessica A. Wachter and Michael Jacob Kahana

July 12, 2019

Summer Institute: Behavioral Macro

Broad Motivation

- Expected utility theory (Savage, 1954) starts with a probability space and an information structure.
- The agent must associate a value with every subset of the space that is consistent with the laws of probability.
- However, it is hard to think about most actual decision problems in this way.

Facts about memory

- The memory system maintains a record of associations between features of the environment and an internal context.
- The features that are present in the environment cue the context, which tells the agent what information is most relevant.
- This idea explains the classic "Laws of Association," which hold across subjects and settings:
  - Recency: better memory for recently experienced items
  - Semantic proximity: better memory for items that have similar meaning
  - Temporal contiguity: better memory for items experienced close in time to a just-remembered item

Related Models

- Mullainathan (2002):
  - The agent recalls an event with greater probability if similar to current events;
  - past recall of an event increases the likelihood of future recall.
- Bordalo, Gennaioli, Shleifer (2019): Physical context cues the agent in ways that are possibly irrelevant.

Features and context

- The agent observes *features* of the environment.
  - We represent features f_t as basis vectors in R^n.
- The agent possesses a *context* c_t, a persistent mental state.
  - The context is a norm-1 vector in R^m.
- Context and features are related via a network of associations.

The network of associations

- Features vector:

$$f_t = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

- Context:

$$c_t = \begin{bmatrix} c_{1t} \\ c_{2t} \end{bmatrix}$$

- The feature-to-context matrix:

$$\begin{aligned} W_t^{f \to c} &= W_{t-1}^{f \to c} + c_t f_t^\top \\ &= W_{t-1}^{f \to c} + \begin{bmatrix} c_{1t} \\ c_{2t} \end{bmatrix} [1, 0] \\ &= W_{t-1}^{f \to c} + \begin{bmatrix} c_{1t} & 0 \\ c_{2t} & 0 \end{bmatrix} \end{aligned}$$

Retrieved context

- The new context is a weighted average of past context and *retrieved context*:

$$c_t = \rho_t c_{t-1} + \zeta c_t^{\text{in}}, \qquad \rho_t \approx 1 - \zeta$$

- Retrieved context:

$$\begin{aligned} c_t^{\text{in}} &= \frac{W_{t-1}^{f\to c} f_t}{||W_{t-1}^{f\to c} f_t||} \\ &\propto \sum_{s=0}^{t} (c_s f_s^\top) f_t \\ &= \sum_{s=0}^{t} c_s (f_s^\top f_t). \end{aligned}$$

- Features call up the context corresponding to when the features were last observed

From memory to decision-making



- In a memory model, context determines the probability of recall, and the speed ("reaction time and percent correct").
- This process is mechanistic.
- In economics, we require a decision-maker.
- We map the memory model into an economic model, by using expected utility, but the probabilities come from memory.

Summary thus far

- Context evolution:

$$c_t = \rho c_{t-1} + \zeta c_t^{\text{in}}, \qquad \rho = 1 - \zeta$$

- Context retrieval:

$$c_t^{\text{in}} \propto W_{t-1}^{f \to c} f_t,$$

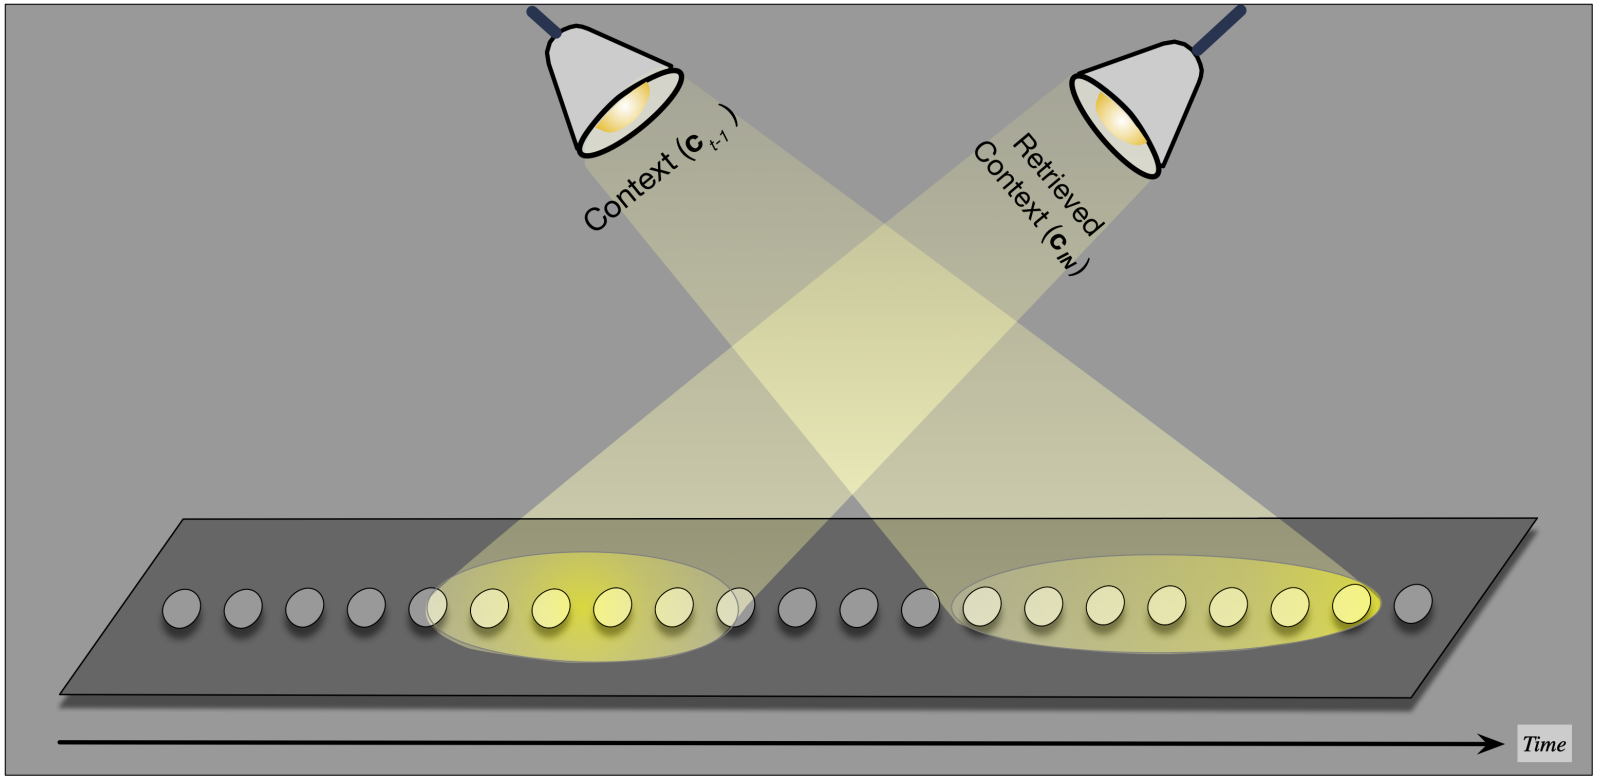
- Features to context matrix:

$$W_t^{f \to c} = \frac{1}{t + \tau} \sum_{s=-\tau}^{t} c_s f_s^\top$$

where τ is the length of the prior sample.

The memory model (cont.)

Think of context as shining a light on previous memories



which we can think of as possible future states of the world.

Outline

Three applications:

1. Over-persistence of memory (Malmendier & Nagel, 2011).
2. Context retrieval and the jump back in time with an application to the financial crisis.
3. The effect of irrelevant external stimuli on decision-making (Guiso, Sapienza, Zingales, 2018).

In paper but not in talk: sticky context and price momentum.

Persistence of memory

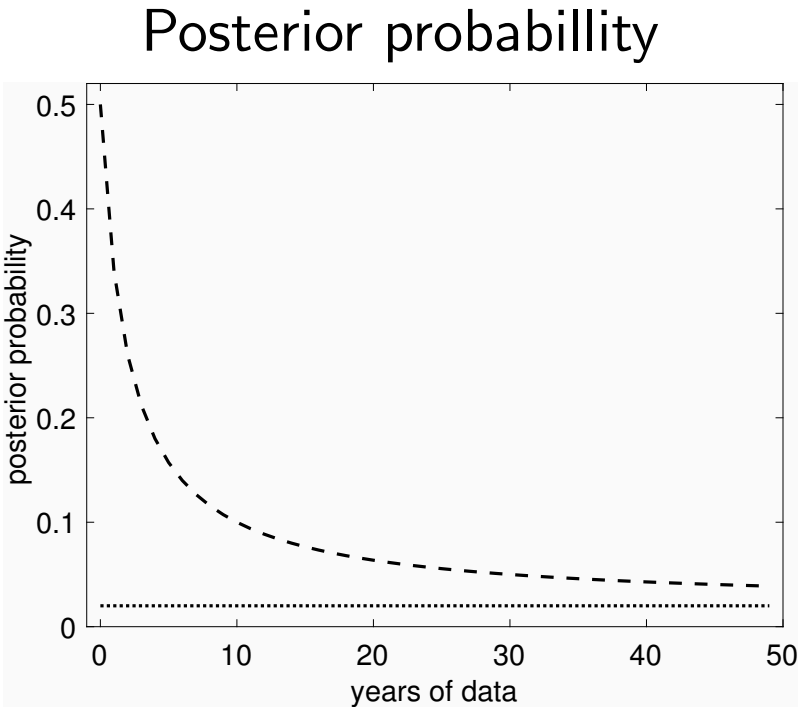
- Excess return states $\tilde{r}(gain) = 1 + \sigma, \tilde{r}(loss) = 1 - \sigma$
- Labor income states: $\tilde{y}(normal) = y > 0, \tilde{y}(depression) = 0$
- Joint contingencies:

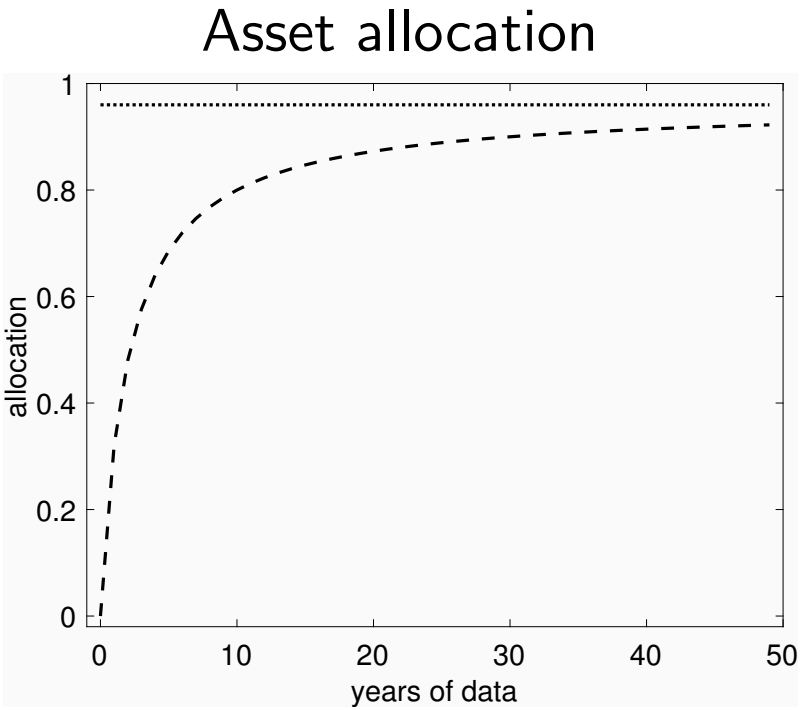
$$\begin{aligned} P &= \begin{bmatrix} P(gain \ \& \ normal) & P(loss \ \& \ normal) \\ P(gain \ \& \ depression) & P(loss \ \& \ depression) \end{bmatrix} \\ &= \begin{bmatrix} \frac{1}{2} & \frac{1}{2} - p \\ 0 & p \end{bmatrix}. \end{aligned}$$

- Risky asset allocation (mean-variance)

$$\pi = \frac{1 - py\sigma}{\sigma^2}$$

Posterior probability and asset allocation: Bayesian case





- The agent starts with a probability of a depression of 1/2. The true probability is 1/50.

Memory for return and labor states

- The stock market return constitutes the features:

$$\text{gain: } f_t = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

$$\text{loss: } f_t = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

- The labor-income state constitutes context.
- In this example, we focus only on the effect of memory on associations through $W^{f \to c}$, setting $\zeta = 1$.

Retrieved Context

- Agent starts with $W_0^{f\to c} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} - p^* \\ 0 & p^* \end{bmatrix}$, $p^* > p$.
- In case of a gain:

$$c_1 \propto W_0^{f\to c} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{2} \\ 0 \end{bmatrix} \Longrightarrow c_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

  The agent recalls the normal labor income state

- In case of a loss:

$$c_1 \propto W_0^{f\to c} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} \frac{1}{2} - p^* \\ p^* \end{bmatrix} \Longrightarrow c_1 = \begin{bmatrix} 1 - 2p^* \\ 2p^* \end{bmatrix}$$

  The agent recalls the depression.

Retrieved Context (cont.)

- Storing the recent event in memory:

$$W_1^{f\to c} = \frac{\tau}{1+\tau} W_0^{f\to c} + \frac{1}{1+\tau} c_1 f_1^\top$$

- In case of a gain:

$$c_1 f_1^\top = \begin{bmatrix} 1 \\ 0 \end{bmatrix} [1, 0] = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}.$$

- In case of a loss:

$$c_s f_s^\top = \begin{bmatrix} 1-2p^* \\ 2p^* \end{bmatrix} [0, 1] = \begin{bmatrix} 0 & 1-2p^* \\ 0 & 2p^* \end{bmatrix}.$$

The depression becomes part of memory, even if it has not occured.

Retrieved context (cont.)

- More generally
  - In case of a gain:

$$c_t \propto W_{t-1}^{f\to c} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \propto \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

  - In case of a loss:

$$c_t \propto W_{t-1}^{f\to c} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \propto \begin{bmatrix} 1-2p^* \\ 2p^* \end{bmatrix}$$

- Therefore

$$\text{plim}_{t\to\infty} W_t^{f\to c} = P^*.$$

- We extract the agents' beliefs about the joint stock return/labor income state from $W_t^{f\to c}$.
- This matrix remains constant at its initial value.

Model of the financial crisis (jump back in time)



1. Standard model of economic disasters
2. The memory model, which turns this into a model of a financial crisis

Standard asset pricing model

- Representative agent model
- Endowment process

$$\log C_{t+1} = \log C_t + \mu + u_{t+1} + v_{t+1}, \tag{1}$$

where u_{t+1} and v_{t+1} are independent, $u_{t+1} \sim N(0, \sigma^2)$ and

$$v_{t+1} = \begin{cases} 0 & \text{with prob. } e^{-p} \\ \log(1-\xi) & \text{with prob. } 1 - e^{-p} \end{cases} \tag{2}$$

- Dividends

$$\log D_{t+1} = \log D_t + \mu + u_{t+1} + \lambda v_{t+1} \tag{3}$$

- Agent

$$E_t \sum_{s=1}^{\infty} \beta^s \log C_s$$

Asset prices

- Intertemporal marginal rate of substitution

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma},$$

- Euler equation:

$$P_t = E_t \left[M_{t+1} (P_{t+1} + D_{t+1}) \right],$$

- Solution:

$$\frac{P_t}{D_t} = \sum_{n=1}^{\infty} \Phi(p)^n = \frac{\Phi(p)}{1 - \Phi(p)}.$$

for

$$\Phi(p) = \beta \left(e^{-p} + (1 - e^{-p}) E \left[(1 - \xi)^{\lambda - \gamma} \right] \right)$$

When $\lambda > 1$, an increase in p (in a comparative statics sense), lowers the price.

Memory model

- Features

$$\text{no crisis: } f_t = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

$$\text{crisis: } f_t = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

- Context

$$c_t = \begin{bmatrix} P(\text{normal}) \\ P(\text{depression}) \end{bmatrix}$$

- Associations

$$W_0^{f \to c} = \begin{bmatrix} P(\text{no crisis \& normal}) & P(\text{crisis \& normal}) \\ P(\text{no crisis \& depression}) & P(\text{crisis \& depression}) \end{bmatrix}$$

$$= \begin{bmatrix} 1 - p^c & p^c(1-q) \\ 0 & p^c q \end{bmatrix}$$

- Assume $p^c = 2.5\%$, $q = \frac{1}{2}$, $\tau = 100$.

Pre-crisis

- The agent experiences neutral features:

$$f_t = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad t = -1, t = -2, \ldots$$

- Retrieved context indicates no crisis:

$$\begin{aligned} c_t^{\text{in}} &\propto W_{t-1}^{f \to c} f_t \\ &\propto \begin{bmatrix} 1 - p^c & p^c(1-q) \\ 0 & p^c q \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} 1 \\ 0 \end{bmatrix} \end{aligned}$$

- Because

$$c_t = \rho c_{t-1} + \zeta c_t^{\text{in}}$$

- The agent places zero weight on depression

$$c_0 \approx \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

Crisis

- Financial crisis

$$f_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

The crisis reinstates the depression context:

$$c_1^{\text{in}} \propto \begin{bmatrix} 1-p^c & p^c(1-q) \\ 0 & p^c q \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \propto \begin{bmatrix} 1-q \\ q \end{bmatrix}$$

- The agent is still a bit in the "old world"

$$\begin{aligned} c_1 &= \rho \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \zeta \begin{bmatrix} 1-q \\ q \end{bmatrix} \\ &= 0.8 \begin{bmatrix} 1 \\ 0 \end{bmatrix} + 0.2 \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} \\ &= \begin{bmatrix} 0.9 \\ 0.1 \end{bmatrix} \end{aligned}$$

- Subjective probability of a depression is 10%

Crisis (cont.)

- Suppose the agent continues to observe crisis features:

$$f_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$c^{\text{in}} \propto W_1^{f \to c} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} .6 \\ .4 \end{bmatrix}$$

$$c_2 = .8 \begin{bmatrix} .9 \\ .1 \end{bmatrix} + .2 \begin{bmatrix} .6 \\ .4 \end{bmatrix}$$

- Probability of a depression is now 16%

The experience of the crisis affects memory

- Memory before the crisis

$$W_0^{f\to c} = \begin{bmatrix} 0.9750 & 0.0125 \\ 0 & 0.0125 \end{bmatrix}$$

- After one crisis observation

$$W_1^{f\to c} = \begin{bmatrix} 0.9653 & 0.0213 \\ 0 & 0.0134 \end{bmatrix}$$

- After two crisis observations:

$$W_2^{f\to c} = \begin{bmatrix} 0.9559 & 0.0293 \\ 0 & 0.0148 \end{bmatrix}$$

Recovery from the crisis

- Suppose neutral features return:

$$c_3^{\text{in}} \propto W_2^{f \to c} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

$$\begin{aligned} c_3 &= \rho \begin{bmatrix} 0.84 \\ 0.16 \end{bmatrix} + \zeta \begin{bmatrix} 1 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} 0.87 \\ 0.13 \end{bmatrix} \end{aligned}$$

Recovery (cont.)

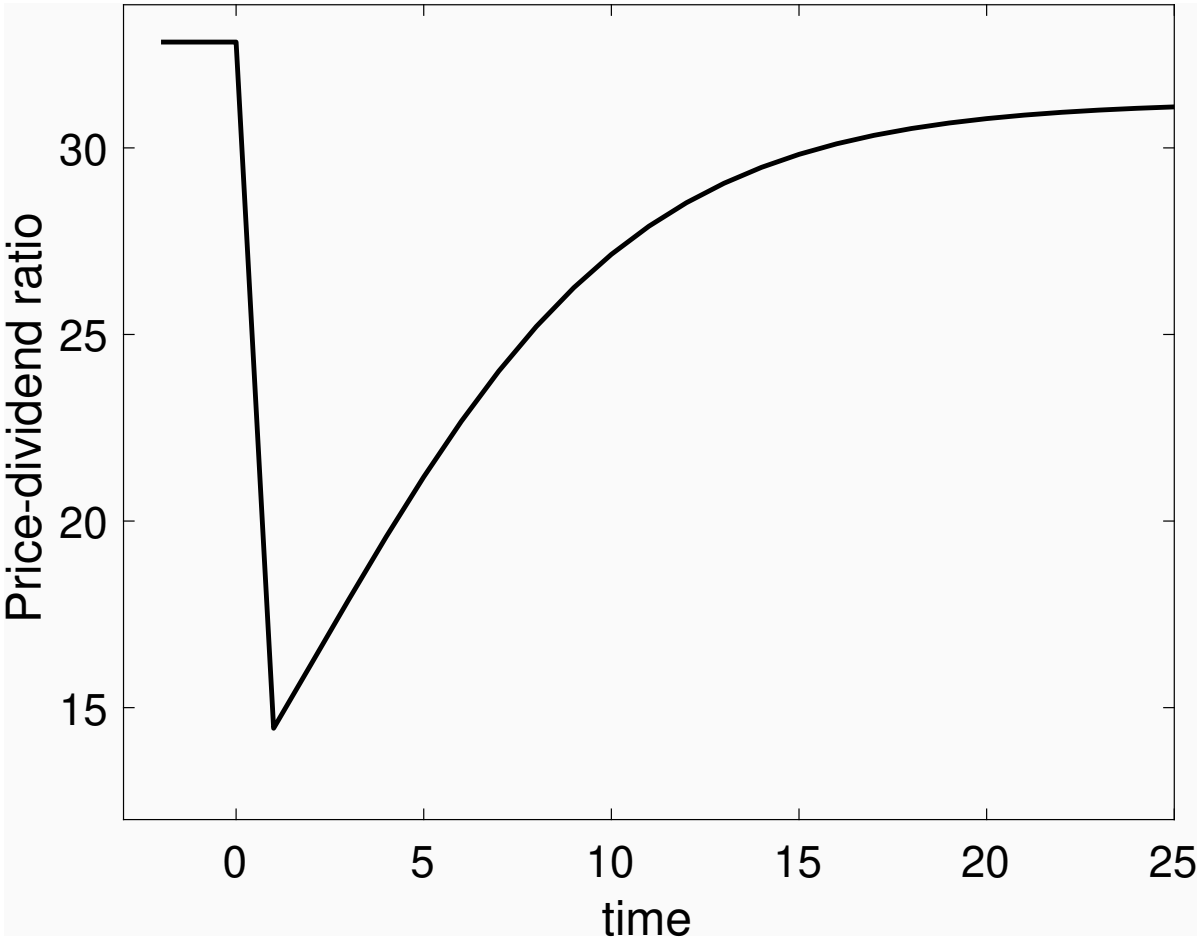
- However, the agent continues to remember the "depression"
- and now the depression mixes with the neutral features

$$W_3^{f\to c} = \begin{bmatrix} 0.9551 & 0.0291 \\ 0.0012 & 0.0146 \end{bmatrix}$$

- Even as the crisis fades from memory, and the agent continues to observe neutral features, the probability of a depression state remains:

$$Wfc_4 = \begin{bmatrix} 0.9546 & 0.0288 \\ 0.0022 & 0.0145 \end{bmatrix}$$

The price-dividend ratio



Summary so far

1. In the first model, context reinstatement leads agents to over-state probability of a depression state
   - This meant overstating the covariance between depression and poor stock market performance.
   - In a quadratic utility model, this implied lower portfolio allocation.
2. In the second model, context reinstatement following a rare event (crisis) led to a jump back in time.
   - The crisis itself was unimportant.
   - However, the appearance of the crisis changed the distribution of stock returns
   - The probability of depression given crisis remains inflated.

Fear and asset allocation

- Guiso et al. (2018) assess changes in risk attitudes toward a lottery under two conditions
  1. Before and after the actual financial crisis (required premium doubles).
  2. Before and after watching a horror movie, for those who dislike horror movies (required premium increases by about 50%).
- In neither case could the distribution of outcomes have changed.

Asset allocation problem

- The agent solves

$$\max_{\pi} E \log(1 + \pi \tilde{r} + \tilde{y}).$$

where

$$\tilde{y} = \begin{cases} 0 & \text{with probability } 1 - p \\ -\xi & \text{with probability } p, \end{cases}$$

with $\xi \in [0, 1]$.

- Assume $\tilde{r}$ also takes on two possible outcomes (each with equal probability), and has mean μ and standard deviation σ.
- Think of $\tilde{y}$ as required expenditures on health, or, say mortgage net of labor income.
- Variation in p leads to variation in risk aversion.

Features

Features space

$$f_t = \begin{cases} e_1 & \text{if} \quad \text{no danger \& no crisis} \\ e_2 & \text{if} \quad \text{danger \& no crisis} \\ e_3 & \text{if} \quad \text{danger \& crisis} \end{cases}$$

where e_j is the jth basis vector in 3 dimensional space

- For danger, think of something causing risk to human capital.
- A financial crisis always represents danger. However, danger need not be financial crisis.

Context space: whether $\tilde{y} = -\xi$ or not ("risk").

Probability space for features and context

Joint contingency matrix

$$P = \begin{bmatrix} \Pr(nr, nd, nc) & \Pr(nr, d, nc) & \Pr(nr, d, c) \\ \Pr(r, nd, nc) & \Pr(r, d, nc) & \Pr(r, c, d) \end{bmatrix}$$

$$= \begin{bmatrix} 1-p & 0 & 0 \\ 0 & p(1-q) & pq \end{bmatrix}$$

Here, nr = no risk, r = risk, etc.

- If there is no risk, there is no danger or crisis.
- If there is risk, then there is definitely danger, which might take the form of a financial crisis.
- Note: stock returns are uncorrelated with $\tilde{y}$, danger, or crisis.

The experiment

- We assume the agent begins with the correct joint probability distribution:

$$W_0^{f\to c} = P$$

- as well as the correct marginal distribution of risk:

$$c_0 = [1-p, p]^\top$$

- The stimulus represents $f_1 = e_2$ (danger, no crisis).

$$c_1^{\text{in}} \propto \begin{bmatrix} 1-p & 0 & 0 \\ 0 & p(1-q) & pq \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \Longrightarrow c_1^{\text{in}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

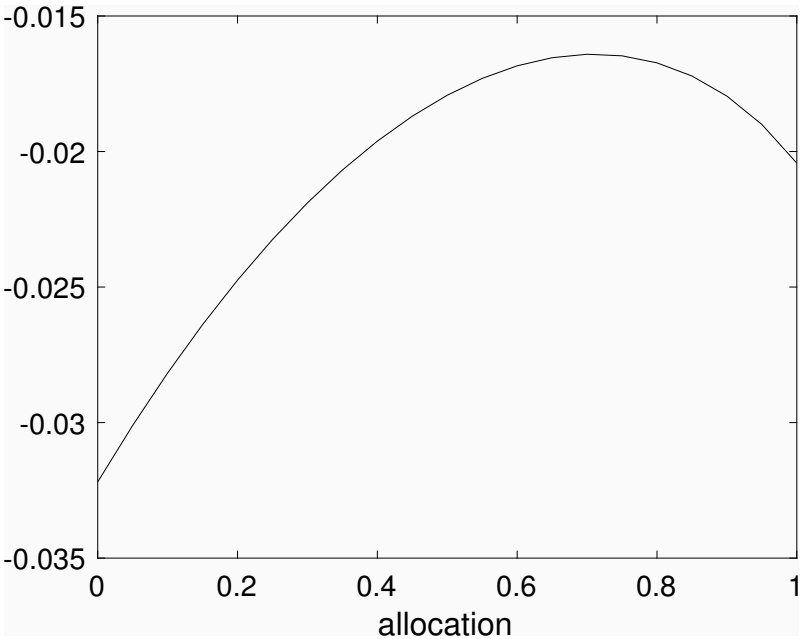
- Therefore, the new context is:

$$c_1 = \rho c_0 + \begin{bmatrix} 0 \\ \zeta \end{bmatrix},$$

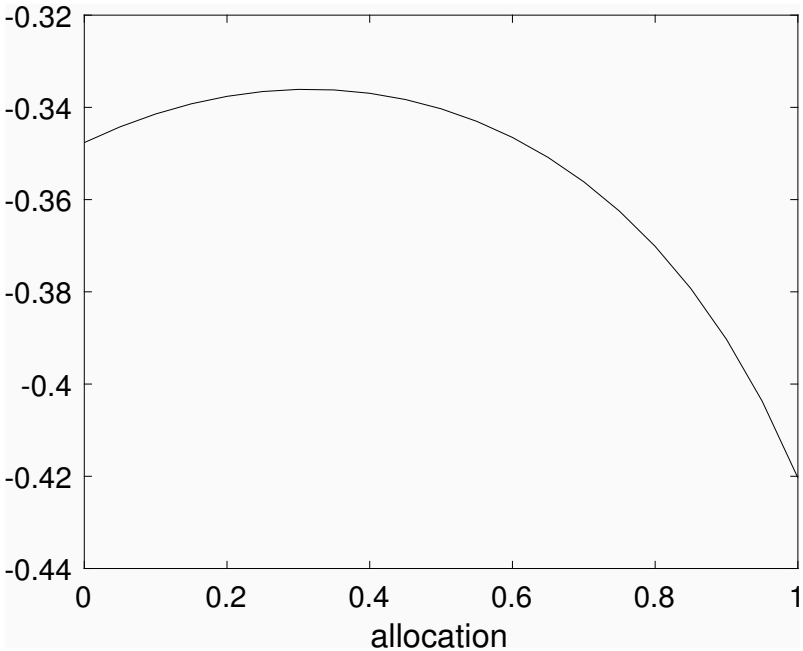
so that the risk probability goes up by 20 percentage points.

Expected utility under context manipulation

Prior beliefs



After context manipulation



Expected utility as a function of the portfolio.

Parameters: $\mu = 4\%$, $\sigma = 20\%$, $p = 2\%$, $\xi = -0.8$

Conclusions

- Our past experiences, and our knowledge about the world, constitute a vast database of information that potentially informs every decision we make.
- Context (endogenous and dynamic) provides a means of retrieving this information when it is most relevant.
- We introduce memory into decision-making by linking context to the beliefs of an economic agent.
- We apply this framework to illustrative problems in financial economics.
- The framework allows for non-Bayesian behavior at two time scales: decisions can be affected by (irrelevant) new information,
- and incorrect probabilities can persist, virtually indefinitely.