



# Collaborative Research Center 649 Workshop

## Risk Preferences and Decisions under Uncertainty

31. January – 1. February 201

*Individual & Contractual Answers to Risks*  
**Prof. Ferdinand Vieider**  
Social Science Research Center Berlin

*Macroeconomic Risks*  
**Prof. Stanley Zin**  
New York University

*Financial Markets and Risk Assessment*  
**Prof. Suleyman Basak**  
London Business School

### Organizers:

Prof. Dr. Wolfgang Härdle  
Prof. Dr. Frank Heinemann  
Prof. Dr. Dorothea Kübler  
Maria Grith

### Cooperation Partners:



**WZB**

Wissenschaftszentrum Berlin  
für Sozialforschung

**Claudia Cerrone** Royal Holloway University of London  
*Education, Disappointment and Public Policy*

**Holger Gerhardt** University of Bonn  
*Cognitive Load Increases Risk Aversion*

**Daniela Grieco** Bocconi University  
*Radical Innovation and Volatility in Stock Prices*

**Maria Grith** HU Berlin *Microeconomic Explanation of the EPK Paradox*

**Christoph Heinzl** French National Institute for Agricultural Research  
*Prudential Saving*

**Jan Heufer** TU Dortmund  
*Nonparametric Comparative Revealed Risk Aversion*

**Paul Karehnke** Paris-Dauphine/Tilburg University  
*Portfolio Choice with Savoring and Disappointment*

**Hong Lan** HU Berlin  
*DSGE Models: Recursive Preferences and Stochastic Volatility*

**Piotr Majer** HU Berlin  
*Risk Patterns and Correlated Brain Activities*

**Baptiste Massenet** University of Lausanne  
*A Theory of Salient Economic Fluctuations*

**Andrija Mihoci** HU Berlin *Cross Country Evidence for the EPK Puzzle*

**Thomas Nitschka** Swiss National Bank  
*Currency Excess Returns & Global Downside Market Risk*

**Marta Serra-Garcia** University of Munich  
*Peer Effects in Risk Taking*

**Yun Shen** TU Berlin  
*Risk-Sensitive Reinforcement Learning for Prospect Theory*

**Christian Stoltenberg** University of Amsterdam  
*Consumption Inequality with Income Inequality*

**Lisette Swart** Tinbergen Institute/University of Amsterdam  
*Irrational Financial Decision-Making?*

**Julian Thimme** University of Muenster  
*Ambiguity in the Cross-Section of Expected Returns*

**Tsvetomira Tsenova** Bulgarian National Bank  
*Monetary policy, Expectations and Uncertainty*

## **Claudia Cerrone**

Royal Holloway University of London

[c.cerrone@lse.ac.uk](mailto:c.cerrone@lse.ac.uk)

*Education, Disappointment and Public Policy* (together with Dan Anderberg)

Individuals undertake investments in education under significant amounts of uncertainty. To account for the risky nature of investments in education, a growing literature on public policy towards education model uncertainty using standard expected utility frameworks. However, it is not clear whether a standard expected utility framework is sufficient to characterize individuals' choices under risk in general and in the context of educational choices in particular. A growing literature in behavioural economics has since the late 1970s documented violations of expected utility and led to a catalogue of theories that generalize expected utility. The central idea of disappointment theory is that an individual forms an expectation about a risky alternative, and may experience disappointment if the outcome eventually obtained falls short of the expectation.

In our paper we introduce the notion of disappointment aversion - modelled, following Gill and Prowse (2012), as loss aversion around the agent's expected payoff where the latter depends on the effort level actually chosen by the agent - into a model of risky investments in education.

The purpose of doing so is two-fold. First, we want to see if disappointment aversion can help shed light on any empirical stylized fact regarding individuals' educational investment behaviours. Second, we want to see how potential disappointment aversion would affect optimal public policy targeted at encouraging participation in education.

In order to achieve our first aim, we explore the impact of disappointment aversion in a simple model of education under risk where agents differ in academic ability and in the absence of policy. In the model, academic ability and individual effort each increases the individual's probability of achieving a high market wage, and are complements in doing so. Using this model we show that disappointment aversion produces a laissez-faire equilibrium where there is a marked negative relationship between the

individual's academic ability (and the level of investment in education) and her marginal rate of return on education. Intuitively this comes about as less academically talented agents are discouraged from increasing their effort as doing so would raise their expectations and hence the disappointment associated with a negative outcome. Conversely, more academically talented agents are encouraged to exert even more effort in order to secure a favourable outcome and avoid disappointment. We discuss this model feature and argue that it is compatible with empirical evidence from set of related literatures: (i) the instrumental variables literature on the rate of return on education, (ii) the literature on the role of credit constraints in college participation, and (iii) the literature that estimates structural models of investments in education.

In order to achieve our second aim, we then introduce government policy into the model in the form of education and income tax policy. We find that disappointment aversion justifies more redistribution (in the form of education subsidies and progressive taxation on income). Indeed, under disappointment aversion, redistribution can have efficiency-enhancing effects by reducing the influence of disappointment considerations on individual education decisions.

## **Holger Gerhardt**

Center for Economics and Neuroscience

Rheinische Friedrich-Wilhelms-Universität Bonn

[holger.gerhardt@uni-bonn.de](mailto:holger.gerhardt@uni-bonn.de)

*Cognitive Load Increases Risk Aversion* (together with Guido Biele, Harald Uhlig, Hauke R. Heekeren)

How people respond to risk is key in many areas of economics. A prominent account of decision making under risk is the dual system approach. It stipulates that decisions are formed by interacting decision-making processes, executed by a fast, risk-averse “emotional” and a slower, less risk-averse “cognitive” system in the brain. Consequently, occupying the cognitive system by another task should increase risk aversion.

We investigated this conjecture experimentally, using a binary lottery choice task and a distractor task. We find that under cognitive load, subjects

were significantly more risk-averse and faster in their lottery choices. Moreover, subjects responded faster when they chose the less risky lottery. All findings are in line with the dual-system approach, but we also discuss how our findings could be explained within a unitary-system framework.

### **Daniela Grieco**

*National Research Council, ISTC & Bocconi University*

[daniela.grieco@unibocconi.it](mailto:daniela.grieco@unibocconi.it)

#### *Radical Innovation and Volatility in Stock Prices*

This paper focuses on the relationship between innovation and volatility of stock market prices, and investigates it through a model depicting innovation as an ambiguous decision. Firms that are more R&D intensive are characterized by a higher degree of idiosyncratic risk: the more radical the innovative process, the stronger the uncertainty of expected future profits. Innovation is a process characterized by Knightian uncertainty (Knight, 1921) or ambiguity, because no distribution of probability can be associated to the success of the investment in R&D. Investing in radical innovation has an uncertain outcome: it creates both favourable expectations for its future growth and fears that the investment will lead to a “dry hole”.

Insofar, scholars have explained a firm’s innovative attitude with its dimension and/or the intensity of market competition (Mazzucato, 2006). However, investing in innovation strongly affects the firm’s stock value: the entrepreneur who wants to push the value of her firm upwards should enhance the firm’s chances of future success, and being innovative is the main way to reach this goal. Some empirical works (Campbell and Shiller, 1988; Mazzucato and Semmler, 1999; Campbell, 2000; Mazzucato, 2003; Mazzucato and Tancioni, 2005; 2012) have emphasized the existence of a link between the degree of innovative disruptiveness and the volatility stock returns at the firm level. At the best of my knowledge, no theoretical studies have been devoted yet to the relationship between innovation and volatility. I introduce a model that aims at summarizing the key mechanisms behind this link and that emphasizes the prominent role of ambiguity in affecting the

decisions related to radical innovation and the consequences of agents’ ambiguity aversion. The model compares ambiguity à la Gilboa-Schmeidler (1989) introducing multiple priors (maximin expected utility) to a smoother form of ambiguity close to the one described in Klibanoff et al.(2005). In the model, the introduction of a radical innovation is captured by allowing the firm to use a radically new input whose cost is sunk and whose returns are ambiguous.

The results show that: (a) firms with the highest R&D intensity face the highest idiosyncratic ambiguity level and have the highest stochastic discount factor; (b) firms introducing radical innovations have higher volatility in their stock returns; (c) the specific form of ambiguity strongly shapes the results.

### **Maria Grith**

Humboldt-University Berlin

[gritmari@hu-berlin.de](mailto:gritmari@hu-berlin.de)

*Microeconomic Explanation of the EPK Paradox* (together with Wolfgang K. Haerdle and Volker Kraetschmer)

Supported by several recent investigations the empirical pricing kernel (EPK) paradox might be considered as a stylized fact. Some authors suggest that this paradox might be caused by regime switching in financial markets. Based on an economic model with state dependent utilities for the financial investors we want to emphasize a microeconomic view that succeeds in explaining the paradox via state dependent preferences. We shall also develop and investigate calibration problems in terms of data fits for basic values of the pricing kernel.

## **Christoph Heinzl**

French National Institute for Agricultural Research

[ch.heinzl@gmail.com](mailto:ch.heinzl@gmail.com)

### *Prudential Saving*

Prudence is a behavioural attitude that is broadly applicable to settings involving risk. It has particular importance in intertemporal choice theory, where it can be interpreted as the intensity of intertemporal substitution. Prior laboratory experiments to elicit prudence have addressed it in a pure-risk sense, by examining behaviour in static lotteries and other gambles.

It is tempting to impute these results into an intertemporal context, leveraging the fact that risk aversion and elasticity of intertemporal substitution are directly mappable under univariate discounted expected utility. However, many empirical studies of intertemporal behaviour suggest that the two ideas may be distinct. To address prudence in its intertemporal sense, we instead design a small-scale laboratory experiment around a two-period consumption/savings model. The utility concept in this model disentangles risk preferences from intertemporal preferences, and suggests the type of exogenous variation to present to subjects in the experiment. The experiment uses a constrained “fill in the blank” design with scenarios involving either income risk or interest-rate risk. In each scenario, subjects must choose how much of their first-period income to save for the second period. The design also implements field-like wealth levels and real time lags to ameliorate the possibility of the decisions being a laboratory artefact. We estimate risk and intertemporal preferences at the individual level using a subject's savings data and the model's structural Euler equation. Excluding outliers, the average coefficient of relative risk aversion is 2.06, the average elasticity of intertemporal substitution is 0.75, and the average coefficient of relative prudence is 3.90. These averages mask a good deal of subject-level heterogeneity, as the respective coefficients of variation are, at a minimum, 70%.

Collaborative Research Center 649 Workshop, 31. January – 1. February 2013  
Risk Preferences and Decisions under Uncertainty

## **Jan Heufer**

TU Dortmund

[jan.heufer@tu-dortmund.de](mailto:jan.heufer@tu-dortmund.de)

### *Nonparametric Comparative Revealed Risk Aversion*

We introduce a nonparametric method to compare risk aversion of different investors based on revealed preference methods. Using Yaari's (1969) definition of “more risk averse than”, we show that it is sufficient to compare the revealed preference relations of two investors. This makes the approach operational; the central rationalisability theorem provides strong support for this approach. The approach is an alternative or complement to parametric approaches and a robustness check. As a necessary first step towards this comparative approach we show how to test data for consistency with stochastic dominance relations, which can also be used to recover larger parts of preferences. We include an application to experimental data by Choi et al. (2007) which demonstrates the potential of the comparative approach.

## **Paul Karehnke**

Université Paris-Dauphine/Tilburg University

[pkarehnke@gmail.com](mailto:pkarehnke@gmail.com)

### *Portfolio Choice with Savoring and Disappointment* (together with Elyès Jouini and Clotilde Napp)

We revisit the model of subjective belief formation proposed by Gollier and Muermann (see Gollier, C. and A. Muermann, 2010, Optimal choice and beliefs with ex-ante savoring and ex-post disappointment, Management Sci., 56, 1272-1284, hereafter GM). In GM, for a given lottery, individuals form anticipated expected payoffs which solve a trade-off between ex-ante savoring and ex-post disappointment. The set of possible anticipated expected payoffs is assumed to be exogenously fixed, independently of the

characteristics of the lottery itself. In particular, in GM, the anticipated expected payoff of a lottery may lie outside the range of objectively possible payoffs. We argue that this is not natural and propose endogenous bounds for the set of possible anticipations equal to the lowest and the highest objectively possible payoff. This permits to compare in a consistent manner lotteries with different supports, to evaluate a lottery without introducing a “framing effect” and to revisit the portfolio choice problem. We obtain new conclusions and interesting insights. In the revisited model, zero-mean risks may be desirable because they weakly increase the range of possible anticipated expected payoffs and may then enable individuals to form optimal anticipated expected payoffs which are more favourable in terms of the savoring and disappointment trade-off. We show that our revisited model is similarly able to explain other empirically observed puzzles like a preference for skewed returns and under-diversification of portfolios.

### **Hong Lan**

Humboldt University Berlin

[lanhongken@gmail.com](mailto:lanhongken@gmail.com)

#### *DSGE Models: Recursive Preferences and Stochastic Volatility*

We analyze the population moments of a nonlinear approximation to a business cycle model with stochastic volatility and recursive preferences. We find that activating a heteroskedastic volatility operationalizes a time-varying risk sensitive channel that alters the patterns of auto and cross correlations of the model while leaving its ability to match several key macro and asset pricing facts untouched. Our method decomposes moments into contributions from shifts in the distribution of future shocks (precautionary effect) and realized shocks (amplification effect), enabling us to isolate the channel through which stochastic volatility operates.

While nonlinear perturbation methods can be applied to solve a DSGE model with recursive preferences and stochastic volatility (see, e.g., Fernandez et al. (2007, 2012)), the population moments of a nonlinearly

approximated perturbation solution are generally not available. Simulating moments is, however, not always feasible or desirable for nonlinear approximations, as such approximations are potentially explosive (Kim et al. (2008)). We use the nonlinear moving average perturbation solution derived in Lan and Meyer-Gohde (2012), to calculate population and simulated moments. With this nonlinear moving average, we are able to decompose population moments past precautionary-amplification components of endogenous variables into the precautionary and amplification contributions of individual shocks. Hence we can more readily identify the contribution of a single shock, by itself and through its interaction with other shocks, to the total variation in a nonlinear environment.

Tallarini (2000), van Binsbergen et al. (2012) and Rudebusch and Swanson (2012) among others use recursive preferences in DSGE models to help reconcile business cycle and asset market facts. Fernandez et al. (2007) and Justiniano and Primiceri (2008) add stochastic volatility to DSGE models with recursive preferences to study the reduction in volatility of U.S. economy since 1980's, as documented in Stock and Watson (2003) and Sims and Zha (2006). Our model, solved by nonlinear moving average perturbation, reproduces Tallarini's (2000) results and captures the contribution of risk to variable with the 3rd order approximation. We show that the expected risk premium and conditional market price of risk operate entirely in the time-varying risk correction channel. Moreover, the addition of stochastic volatility does not contribute to the ability of the model to match the unconditional market price of risk when the volatility of log consumption growth is held constant. Adding stochastic volatility, however, enables the model to produce a conditional market price of risk and an expected risk premium with significant time variation.

To sum up, we have shown the effects of introducing stochastic volatility into a DSGE model with recursive preferences on both macro and financial variables by analyzing the simulated and population moments. In order to disentangle the role of stochastic volatility from other shocks and parameters, we propose a closed-form population moments calculation and decomposition based on the nonlinear moving average perturbation method that can be applied to a class of nonlinear DSGE models in the literature

**Piotr Majer**

Humboldt University Berlin

[majerpio@cms.hu-berlin.de](mailto:majerpio@cms.hu-berlin.de)

*Risk Patterns and Correlated Brain Activities* (Alena Mysickova, Song Song, Peter N.C. Mohr, Hauke R. Heekeren, Wolfgang K. Haerdle)

Decision making usually involves uncertainty and risk. Understanding which parts of the human brain are activated during decisions under risk and which neural processes underlie (risky) investment decisions are important goals in neuroeconomics. Here, we reanalyze functional magnetic resonance imaging (fMRI) data on 17 subjects which were exposed to an investment decision task from Mohr et al. (2010). We obtain a time series of three-dimensional images of the blood-oxygen-level dependent (BOLD) fMRI signals. Our goal is to capture the dynamic behaviour of specific brain regions of all subjects in this high-dimensional time series data, by a flexible factor approach resulting in a low dimensional representation. We apply a panel version of the dynamic semiparametric factor model (DSFM) presented in Park et al. (2009) and identify task-related activations in space and dynamics in time. Further, we classify the risk attitudes of all subjects based on the estimated low-dimensional time series. Our classification analysis successfully confirms the estimated risk attitudes derived directly from subjects' decision behaviour.

**Baptiste Massenet**

University of Lausanne

[baptistemassenot@gmail.com](mailto:baptistemassenot@gmail.com)

*A Theory of Salient Economic Fluctuations* (together with Kenza Benhima)

Human perception tends to be biased towards elements that are unusual or that stand out. Whether we perceive light, weights, or lottery payoffs, our perception is distorted when one object is more salient. For example, a lottery that gives the chance to earn a million dollars, individuals

may tend to overweight the possibility of earning this million and even otherwise risk-averse individuals may end up buying the lottery ticket despite negative expected value. In our paper we study the implications of such a perception bias on the investment behaviour of entrepreneurs. We find that once taken into account, the overweighting of salient payoffs can lead to endogenous cycles of investment. More specifically, entrepreneurs decide how much to invest in a risky project. When the upside of the project is more salient, entrepreneurs tend to overestimate the expected value of the project. This implies that they hire too many workers and the expected value of their project turns negative. By contrast, when the downside is more salient, they tend to underestimate the expected value, do not hire enough workers, and end up making a positive profit.

When studying the dynamics, we find that the economy can exhibit endogenous cycles. When computing the salience of payoffs, we assume that the previous gains or losses of entrepreneurs are reflected in the payoffs considered by entrepreneurs.

After a gain, entrepreneurs will integrate these gains to the possible payoffs, which will make them more risk-seeking. After a loss, we assume that entrepreneurs do not integrate them to the payoffs as in Thaler and Johnson (1990). For example, when considering the same payoffs from the project, the previous positive profit is going to be integrated to the payoffs of the project. When perceiving these payoffs, the upside is going to become more salient, and the entrepreneur is going to overestimate the expected value of his project and make a negative profit. This loss is going to make him more risk-averse next period.

The paper thus offers a theory of economic fluctuations that is not based on economic fundamentals but on the perception of entrepreneurs. This is reminiscent of the animal spirits hypothesis put forward by Keynes. Furthermore, this may help reconcile movements in economic activity that does not seem to be in line with economic fundamentals.

When this economy faces an innovation, for example a positive productivity shock, the movements described above can be amplified. A higher productivity indeed makes the upside of a project more salient and makes entrepreneurs more likely to overinvest.

Importantly, the economy can experience recessions without resorting to negative productivity shocks as in the RBC paradigm. Furthermore, unlike in the news shock literature where recessions occur following overoptimistic signals, recessions can go below the trend.

Finally, we believe that one important advantage of our approach is that recessions are endogenous. By this, we mean that the likelihood of a recession at date  $t + 1$  can be assessed using information from date  $t$ . Such a label is also used by the literature on indeterminacy in business cycle models. However, in this literature, cycles occur following a stochastic process referred to as sunspots. Although like in our work, economic fluctuations do not reflect fundamentals, they are essentially exogenous, or at least unpredictable. By endogenizing economic fluctuations, we believe our approach can be thought of as a first step towards answering the question of how to prevent recessions.

### **Andrija Mihoci**

Humboldt University Berlin  
[mihociax@cms.hu-berlin.de](mailto:mihociax@cms.hu-berlin.de)

*Cross Country Evidence for the EPK Puzzle* (together with Wolfgang K. Haerdle and Maria Grith)

A worldwide evidence for the Empirical Pricing Kernel (EPK) paradox is presented. The pricing kernel is derived under a modified utility function framework, i.e., we consider a state-dependent utility approach. The methodology is applied to portfolios of blue-chip stocks at six worldwide largest stock markets. Based on stock market data we show that statistically there exists a EPK paradox. The results are quite persistent across countries and across the underlying framework specifications. Moreover, the estimated EPK preference parameters exhibit a time-varying pattern.

### **Thomas Nitschka**

Swiss National Bank

[thomas.nitschka@snb.ch](mailto:thomas.nitschka@snb.ch)

*Currency Excess Returns and Global Downside Market Risk*

The difference between current forward and spot exchange rates, i.e. the forward discount, should be a good predictor of future exchange rate movements. However, a wealth of studies initiated by Tryon (1979), Hansen and Hodrick (1980) and Fama (1984) find that exchange rate changes do not follow forward discounts or, equivalently, interest rate differentials. The uncovered interest rate parity (UIP) fails to hold empirically. The presence of a risk premium demanded for foreign currency investments is a potential explanation for UIP violations. Such risk premia may reflect crash risk or rare events (e.g. Brunnermeier et al., 2009; Farhi et al., 2009; Farhi and Garbaix, 2011), peso problems (e.g. Burnside et al., 2011), or differences in the sensitivity of currencies to systematic risk factors (e.g. Ang and Chen, 2010; Christiansen et al., 2011; Lustig and Verdelhan, 2006, 2007; Lustig et al., 2011; Menkhoff et al., 2012a; Verdelhan, 2010, 2012). Especially this latter strand of the literature attracts a lot of attention because it adopts well established theoretical frameworks and empirical methods from time series and cross-sectional analysis of stock and bond returns to the exchange rate context.

However, this literature faces criticisms based on three main grounds. First, it typically relies on formation of currency portfolios to show empirically that there is a relation between systematic risk factors and deviations from the UIP. Second, general pricing models that explain excess returns on currency portfolios deliver implausible values of preference parameters. For example, the consumption-based models examined in Lustig and Verdelhan (2007) explain the cross-sectional dispersion in currency portfolio returns only with implied risk aversion parameter values around 100. A similar value of risk aversion is needed in an intertemporal CAPM setting to explain currency portfolio returns (Galsband and Nitschka, 2012). Third, confronted with bilateral currency excess returns, the most successful

empirical currency pricing models employ rather ad hoc risk factors extracted from currency return data. Examples include the “carry trade factor”, i.e. the return differential on high and low forward discount sorted currency portfolios, introduced by Lustig et al. (2011) or a measure of global volatility on foreign exchange markets (Menkhoff et al., 2012a). More general proxies of systematic risk cannot explain bilateral currency excess returns well (Burnside et al., 2011; Burnside, 2011).

We propose a simple extension of the standard empirical version of the CAPM which implicitly assumes that investor’s utility is logarithmic such that the relative risk aversion parameter is restricted to be close to one. Our preferred empirical model has two additional key features. First, it takes into account country-specific and global market risks from a national investor’s point of view. Second, it distinguishes between global upside and downside risks and allows thus for loss aversion in individual preferences. We define global upside risk as sensitivity to the global component of the market excess return when it is positive. Global downside risk is accordingly measured by the sensitivity to the global component of the market excess return when it is negative. For the sample period from January 1999 to March 2012, this empirical model explains more than 40% of the cross-sectional dispersion in 20 monthly bilateral currency excess returns from a U.S. investor’s perspective. The distinction between country-specific and global risk combined with the distinction between global upside and global downside risk is crucial for explaining the average currency excess returns under study. We find that sensitivity to the global downside risk component of the market return is significantly priced. We also show that downside risk is economically important. In line with estimates from downside risk models confronted with firm-level stock returns (Ang et al., 2006), a two standard deviation increase in the exposure to global downside risk would lead to an 8% p.a. increase in currency excess returns. Hence, an empirical model that restricts risk aversion to low values but allows for loss aversion helps to explain the cross-sectional differences in bilateral currency excess returns. Finally, we highlight that downside risk is reflected in the carry trade factor introduced by Lustig et al. (2011).

**Marta Serra-Garcia**

University of Munich

[Marta.SerraGarcia@lmu.de](mailto:Marta.SerraGarcia@lmu.de)

*Peer Effects in Risk Taking* (with Amre Lahno)

This paper examines the effect of peers on individual risk taking. In the absence of informational motives, we investigate why social utility concerns may drive peer effects. We test for two main channels: utility from payoff differences and from conforming to the peer. We show experimentally that social utility generates substantial peer effects in risk taking. These are mainly explained by utility from payoff differences, in line with outcome-based social preferences. Contrary to standard assumptions, we show that estimated social preference parameters change significantly when peers make active choices, compared to when lotteries are randomly assigned to them.

**Yun Shen**

TU Berlin

[yun@ni.tu-berlin.de](mailto:yun@ni.tu-berlin.de)

*Risk-Sensitive Reinforcement Learning for Prospect Theory* (with Wendelin Boehmer, Tobias Sommer, Michael Tobia, Christian Buechel, Klaus Obermayer)

Our work aims at modelling risk-sensitive decision-making procedures of human beings in an uncertain environment by incorporating prospect theory (PT) into the framework of reinforcement learning (RL).

Reinforcement learning is widely applied to model human decision-making procedures and similar computational structures have been found in primate brains (Schultz et al., 1997). RL models how agents ought to take actions in an uncertain environment where event probabilities are unknown so as to maximize some notion of cumulative rewards. However, the hypothesis of expected utility, on which RL relies, is continually challenged



by behavioural economic theories like prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). Although the prospect theory fits the human behaviours in descriptive tasks better than the expected utility theory, it is difficult to apply PT in learning tasks that are usually modelled by RL, since in learning tasks, the event probabilities required in the calculation of prospects are unknown for subjects.

In our work, we apply a special type of functions, called utility-based shortfall (USF), which is originally introduced to measure financial risk (Föllmer and Schied, 2004), to replace the calculation of prospects in PT. We show that key effects of PT, e.g., the S-shape subjective probability curves and varied reference points, can be well replicated by setting proper utility functions in USF. We therefore propose a general framework of risk-sensitive Markov decision processes and solve the induced optimization problems by dynamic programming. Using USF, we derive a generalized RL algorithm and prove its convergence under proper conditions. In particular, utility functions in the new algorithm work upon temporal differences (TD), which measure the difference between the expected and actual reward, rather than the actual reward.

The generalized RL algorithm is applied to analyze the data collected from 80 subjects in a 3-decision sequential investment task. Results show that the new algorithm fits the behavioural data significantly better than standard RL algorithms. Based on the best fitted utility functions, we discover that in the task human subjects are in general risk-seeking at the “loss” side while risk-averse at the “gain” side, which is consistent with the prospect theory. We can furthermore categorize subjects into different groups according to their utility functions. The categorization coincides with subject’s performance measured by the total amount of money obtained in the task. Finally, an analysis of fMRI data reveals two significant correlations: i) between TD sequences generated by the new algorithm and activations in ventral striatum; and ii) between Q-value sequences of the new algorithm and activation in medial orbitofrontal cortex.

## **Christian Stoltenberg**

University of Amsterdam

[C.A.Stoltenberg@uva.nl](mailto:C.A.Stoltenberg@uva.nl)

### *Consumption Inequality with Income Inequality (with Vadym Lepetyuk)*

The recent rise in consumption inequality in response to the increase in income inequality in the U.S. is puzzling to expected-utility-based incomplete market models. We show that a model with two-sided lack of commitment and chance attitudes like optimism and pessimism, as emphasized by prospect theory, can explain the relationship and can avoid the systematic bias of the expected utility models. For realistic values of risk aversion and for chance attitudes, the incentives for households to share the idiosyncratic risk decrease. The latter effect endogenously amplifies the increase in consumption inequality, thereby improving the fit to the data.

## **Lisette Swart**

Tinbergen Institute/University of Amsterdam

[l.swart@vu.nl](mailto:l.swart@vu.nl)

### *Irrational Financial Decision-Making?*

Time-inconsistent behaviour of farmers in developing countries has puzzled researchers as it implies foregoing highly profitable opportunities: farmers keep postponing profitable investments which they were determined to do. This behaviour violates the common assumption of exponential discounting and the reasons for this violation are not yet well understood. Many researchers have pointed at self-control problems, but previous studies do not yet yield clear-cut evidence in favour of this explanation. Recently, Gine, Goldberg, Silverman and Yang (2012) investigated alternative causes of people's change of plan, including the effects of social interaction and of shocks. Particularly in rural areas in developing countries, there is a strong communal sense in which community members assist each other financially

when any of them faces a shock. Experiencing such a shock can in turn influence one's optimal path of consumption.

Gine et al. carried out a field experiment to investigate whether shocks determine plan revisions but did not find evidence for it. First, participants' time preferences were elicited in an experiment using a method labelled convex time budget (CTB) by Andreoni and Sprenger (2012a) where stakes equalled a month's salary. In half of the offers, participants had to allocate an amount of money between 'tomorrow' and 'in 1 month,' representing the near time frame, and in the other half between 'in 2 months' and 'in 3 months,' the far time frame. Within each time frame allocating money to the later period yielded some interest and the experiment involved a number of offers with different rates of return. In each household, one of the two spouses was randomly selected to be paid out one of their choices. Subsequently, those participants whose selected offer was in the far period were revisited some days before the disbursement of the first payment (almost 2 months later) and were given an opportunity to revise their initial decision at the same rate of return. Importantly, during this revision the old allocation was clearly shown to the respondent. They found that shocks that people experienced and measures of social assistance did not explain peoples plan revisions.

This paper provides deeper insight into this issue by repeating Gine et al.'s experiment and improving the methodology in three important aspects. First, by using financial diaries weekly data on financial transactions as well as social and health events that people were collected and in this way we were able to measure shocks much more precisely. Second, the equation that is estimated in this paper is directly derived from utility maximization and this derivation shows that Gine et al.'s estimating equation is wrongly specified. Finally, in this paper the boundary conditions implicit in the experiment are taken into account. In the allocations that people have to decide on, consumption that is allocated to any period is restricted to be non-negative and not greater than amounts implied by the experimental budget. As a result, the allocations selected do not represent people's optimal levels of consumption if it involves a corner allocation and it is important to take this into account in the analysis. Contrary to Gine et al., we find a large effect of shocks on plan revisions.

**Julian Thimme**

Westfälische Wilhelms-Universität Münster

[01juth@wiwi.uni-muenster.de](mailto:01juth@wiwi.uni-muenster.de)

*Ambiguity in the Cross-Section of Expected Returns* (with Clemens Voelkert)

This paper investigates if ambiguity aversion is present in investors' decision patterns by looking at the cross-section of stock returns and macroeconomic variables. We estimate a set of preference parameters assuming that investors act in line with the smooth ambiguity model of preference as developed by Klibanoff, Marinacci, and Mukerji. This model generalizes the recursive preference model of Epstein and Zin which we use as a benchmark. The impact of a decision maker's ambiguity attitude on asset prices is characterized by an additional uncertainty factor, i.e. an additional term in the pricing kernel, which is the expected continuation value of future consumption, conditional on the economic model at hand. We introduce a novel methodology for the estimation of this term.

Our point estimates of the ambiguity aversion parameter are between 25 and 40, whereas our risk aversion estimates are considerably lower. The substantial difference indicates that market participants are ambiguity averse. However, due to the small post war sample, a Wald test does not reject ambiguity neutrality. To put this into perspective, we investigate the finite sample behaviour of the GMM estimator in a simulation study. We simulate asset returns with the help of a long run risks model, similar to Bansal and Yaron. We find that even in an artificial setting, in which ambiguity aversion has a perceptible impact on asset prices, ambiguity neutrality cannot be rejected in small samples. An estimation of the preference model of Epstein and Zin, based on asset returns simulated under ambiguity aversion, shows that the risk aversion parameter is estimated with an upward bias, while the cross-sectional pricing performance of this model is only slightly lower than that of the smooth ambiguity model. Based on returns on 30 test portfolios, we evaluate if ambiguity aversion helps explaining the cross-section of expected returns. In line with the findings from the simulation study, we find that incorporating ambiguity attitudes into the decision model slightly

improves the fit to the data while keeping relative risk aversion at more reasonable levels. This can be explained by the positive covariation in the time series of the additional uncertainty factor and the Epstein Zin-risk factors. Neglecting investors' ambiguity aversion leads to an amplification of the risk aversion parameter, which might explain the large values reported in the literature. We argue that the smooth ambiguity model offers a more reasonable explanation of asset prices.

**Tsvetomira Tsenova**

Bulgarian National Bank

[Tsenova.Ts@bnbank.org](mailto:Tsenova.Ts@bnbank.org)

*Monetary policy, Expectations and Uncertainty*

This paper shows that monetary policy does and should respond systematically to time variation in ex-ante uncertainty and heterogeneity in the private sector's forward views over the business cycle. Furthermore, the findings suggest that there is a non-trivial difference between the two factors: while higher uncertainty justifies aggressive monetary policy accommodation, heterogeneity acts in the opposite direction. The study uses as evidence own estimates of individual forecast uncertainty and heterogeneity from professional forecasters' reports, following methodologies developed by D'Amico and Orphanides (2006), Giordani and Söderlind (2003) and Zarnowitz and Lambros (1987). Empirical tests are initially conducted on the basis of an augmented forward-looking Taylor rule framework, modified to account for learning and robustness. Normative explanation is further provided by evaluating the optimal forecast-based monetary policy response under imperfect knowledge, given a set of heterogeneous nested reference structural models, estimated to best-fit private sector's forecasts in addition to contemporaneous data. The analysis incorporates the period of the recent financial and economic crisis.