

Viefers, Paul

Working Paper

Should I stay or should I go? A laboratory analysis of investment opportunities under ambiguity

DIW Discussion Papers, No. 1228

Provided in Cooperation with:

German Institute for Economic Research (DIW Berlin)

Suggested Citation: Viefers, Paul (2012) : Should I stay or should I go? A laboratory analysis of investment opportunities under ambiguity, DIW Discussion Papers, No. 1228, Deutsches Institut für Wirtschaftsforschung (DIW), Berlin

This Version is available at:

<https://hdl.handle.net/10419/61311>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Discussion Papers

Deutsches Institut für Wirtschaftsforschung

2012

Should I Stay or Should I Go? A Laboratory Analysis of Investment Opportunities under Ambiguity

Paul Viefers

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2012

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
<http://www.diw.de>

ISSN print edition 1433-0210
ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
<http://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
<http://ideas.repec.org/s/diw/diwwpp.html>
<http://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

Should I stay or should I go?

A laboratory analysis of investment opportunities under ambiguity

Paul Viefers*

July 26, 2012

Abstract

This paper investigates the impact of uncertainty on an irreversible investment decisions in the laboratory. Subjects own the option to seize a claim on the future sum of realizations from an (ambiguous) random walk. I contrast model predictions of the Subjective Expected Utility model (SEU, Savage, 1954) with model predictions made by Multiple-prior Expected Utility models (MEU, Gilboa & Schmeidler, 1989; Epstein & Schneider, 2003b). Observed investment behavior is at odds with the SEU prediction and deviates in a direction predicted by MEU models. On average, subjects in a treatment group, facing an ambiguous random walk, exhibit an ambiguity premium that presents a mark-up on average reservation profits in a control group. Hence, subjects shun to expose themselves to an ambiguous payoff process and invest later than participants facing a risky payoff process.

JEL-Classification: D08, D83

Keywords: Ambiguity Aversion, Multiple Priors, Optimal Stopping, Irreversible Investment.

*DIW Graduate Center, Mohrenstr. 58, D-10117 Berlin (email: pviefers@diw.de). I would like to thank Georg Weizsäcker, Philipp Strack, John Ham, Jörg Breitung, Tobias Schmidt, and seminar participants at WZB Berlin for helpful discussions. I would also like to thank Jörn Hees (DFKI, Kaiserslautern) for programming the excellent experimental software. All remaining errors are mine.

1 Introduction

Timing matters.

Many everyday decisions are sequential timing or Optimal Stopping decisions. For example, whether to take a job offer or decline it in favor of future offers or whether to invest your savings into financial markets now or later. On a corporate level, firms may want to engage R&D of a new technology not before market situations are 'ripe', or firms plan to time stock markets to maximize the proceeds from an IPO. In general in an Optimal Stopping problem, a decision maker faces a sequence of payoffs (π_1, π_2, \dots) , drawn from some distribution. After each realization of the payoff process, the decision maker has to decide whether to continue sampling or to stop sampling and exercise his option. The payoff $f_\tau(\cdot)$ from stopping may either be a function of past realizations of the process $\{\pi_t\}_{t=1}^\tau$ or future realizations $\{\pi_t\}_{t=\tau}^T$ where T may be infinite. Both such timing problems are Optimal Stopping problems, where in the latter case one often speaks of an Irreversible Investment problem. Due to the fact that they relate to a great number of economic decision problems, Optimal Stopping problems are of considerable interest from both a normative and descriptive viewpoint. This paper takes the descriptive viewpoint and contrasts observed behavior with normative benchmarks.

The theoretical literature on Optimal Stopping in Economics (McDonald & Siegel, 1986; Chow, Robbins & Siegmund, 1971; Dixit & Pindyck, 1994, are classical references) employs tools from stochastic control theory (Snell, 1952) to find the optimal strategy for a decision maker. It turns out that for many classical problems, the optimal strategy is easily determined by the Principle of Optimality: weigh the value from stopping today against the value of continuing until tomorrow, given you behave optimally tomorrow.

What is common to all classical approaches to optimal stopping, is the assumption that there is only one distribution for the payoff process and this distribution is perfectly known to the decision maker. Going from the level of a stock-listed multinational corporation to the level of the individual, the latter part of this assumptions seems more and more problematic.¹ It is well-known since Ellsberg (1961), that in a static choice between ambiguous and risky lotteries, the distinction between risk and uncertainty (in the sense of Knight, 1921) is behaviorally important. This finding was replicated in many subsequent studies (e.g. Halevy, 2007). Based on a thought-provoking paper by Al-Najjar & Weinstein (2009), however, reasonable doubt has been cast on the normative viability of ambiguity models. In this context, Dominiak, Dürsch & Lefort (2009), for example, find experimental evidence for a key point raised by Al-Najjar & Weinstein. They show that in a dynamic

¹In the spirit of Epstein & Schneider (2003a, 2007), Knightian uncertainty may stem from complexity. The financial crisis has highlighted that even companies such as GoldmanSachs find it difficult to fully resolve any uncertainty surrounding certain complex securities. Arora, Barak, Brunnermeier & Ge (2011) provide a rigorous argument that illustrates how limited computational power (a tiny deviation from full rationality) leads to model uncertainty that cannot be resolved.

Ellsberg experiment, the majority of subjects are not time-consistent. A crucial assumption underlying most dynamic ambiguity models.

Against this backdrop, it seems expedient to analyze the predictive accuracy of ambiguity models in various settings. This paper puts the focus on irreversible investment decisions. Such decisions involve the option to seize an investment opportunity with time-varying returns at certain cost.

The results from this experiment indicate that there exists an effect of ambiguity in an Optimal Stopping problem. Subjects in a treatment group, facing an ambiguous payoff process, invest, on average, later than subjects facing a risky payoff process. Given the experimental design, the SEU model predicts exactly the opposite: subjects facing ambiguity should in principle stop no later than subjects facing risk. Intuitively speaking, the results indicate that subjects shun to expose themselves to an ambiguous payoff process and therefore ask for a relatively higher reservation profit to invest in it. This is in line with the predictions made by the MEU model.

The results obtained here confirm results obtained earlier by Asano, Okudaira & Sasaki (2011) and Della Seta, Gryglewicz & Kort (2012). As noted by Miao & Wang (2011), however, it is important to distinguish two fundamentally different decision problems. Those where upon stopping payoffs are ex post certain, versus those where upon investment payoffs are ex post uncertain. While the job search problem falls into the former category, the Irreversible Investment problem falls into the latter. Even though related experimental studies acknowledge this fact (e.g. Della Seta et al., 2012), none carries this distinction into the underlying experimental design. Arguably most investment decisions valid for the individual private investor, seem to be such that payoffs are ex post uncertain. This is not merely a semantic issue. In fact, model predictions about the direction of an ambiguity effect easily reverse. To the author's best knowledge, this paper is the first to explicitly design an irreversible investment experiment in a way that payoffs are *not* ex post certain.

The paper is organized as follows. Section 2 briefly reviews the related literature. Section 3 motivates the underlying theoretical model by a simple example. Section 4 provides the general model setup and provides the central proposition to be tested. Section 5 describes the experimental setup and implementation in the laboratory. In section 6 I conduct a statistical analysis of the data obtained in the laboratory. Section 7 finally concludes.

2 Related literature

Based on the Ellsberg paradox, a great number of models have been developed to rationalize these static choices made by individuals under ambiguity (see *inter alia* Camerer & Weber, 1992; Mukerji & Tallon, 2004; Gilboa, Postlewaite & Schmeidler, 2008; Etner, Jeleva &

Tallon, 2009; Guidolin & Rinaldi, 2010; Epstein & Schneider, 2010, for extensive reviews). A notable feature of this theoretical literature is that there are many competing models of ambiguity, some of which are mutually exclusive and vastly different in certain model predictions. However, not only have the static Ellsberg choices been rationalized by ambiguity models of choice, but many such models were subsequently taken to a dynamic setting (see e.g. Epstein & Schneider, 2010, for a review). There is less evidence on how and whether ambiguity aversion carries over to dynamic contexts.

Theoretical foundations for optimal stopping theory under ambiguous payoff processes are mainly due to Epstein & Schneider (2003b); Nishimura & Ozaki (2007); Riedel (2009, 2010).² These authors take the Multiple-prior Expected Utility (MEU) model of Gilboa & Schmeidler (1989) as the basic building block and suitably adapt it to a dynamic context. Despite the widespread application of optimal stopping models in economic theory, there are relatively few papers that analyze their descriptive accuracy under risk or ambiguity. One reason might be that every empirical analysis with field data will be inevitably marred by many potential confounds. Experiments provide a way to safely control for such.

The largest body of literature on optimal stopping tasks in the laboratory comes from the field of sequential search. For example, Amnon Rapoport and co-authors in a series of papers, investigated the individual performance of subjects in an optimal stopping task under risk experimentally (Rapoport & Tversky, 1966; Kahan, Rapoport & Jones, 1967; Rapoport & Tversky, 1970; Seale & Rapoport, 1997). Similarly, there is a branch of experimental literature putting a focus on testing job search models in the laboratory, e.g. Schotter & Braunstein (1981); Cox & Oaxaca (1989, 2000). More recently, Oprea, Friedman & Anderson (2009) adapt an experimental design replicating the theoretical environment motivated by Dixit & Pindyck (1994). In their paper, subjects face a risky random walk and have the option to earn its current value or forego it in favor of future values. They find that subjects approximate the optimal strategy surprisingly well, but that they tend to stop prematurely.

A first paper making a step towards checking the effect of uncertainty on decisions in the laboratory is given by Cox & Oaxaca (2000). In their experiment, however, participants are endowed with an objective prior over states of the world. Asano et al. (2011) investigate the descriptive power of the job search model by Nishimura & Ozaki (2004) when subjects have no prior information about the distribution of states of the world. Asano et al. (2011) find that observed choices support Nishimura & Ozaki's model in the sense that subjects are willing to accept lower wage offers under ambiguity than under risk. Similarly, Della Seta et al. (2012) analyze the a related situation where subjects do not face a random, but deterministic payoff process and have the option to earn its current value. This option

²Earlier foundations are to be found in the mathematical literature on coherent risk measures (see e.g. Artzner, Delbaen, Eber & Heath, 1999; Riedel, 2004; Föllmer & Schied, 2004) or robust control theory in macroeconomics (Hansen & Sargent, 2001).

may cease before execution, however, leaving subjects with zero payoff. The probability that the investment opportunity expires is ambiguous in their setting. They find that subjects react to ambiguity by exercising the option earlier relative to a control group, as predicted by their model.

3 Motivation through a simple example

In this section, I illustrate the behavioral intuition behind the impact of uncertainty and uncertainty aversion on the decision to invest. Toward that end, I shall take the simple two-period two-state example from Nishimura & Ozaki (2007).

To make exposition viable, I will make a few simplifying assumptions. First, assume a risk-neutral decision maker, facing an irreversible investment opportunity. Second, let the one-period discount rate δ be equal zero. Third, let there be only two periods $t = 0, 1$ where any uncertainty has already been resolved in period 0. The decision maker contemplates whether to invest in period $t = 0$ in period $t = 1$ or not at all. In order to seize the investment opportunity, investment costs I have to be incurred. The immediate period-0 profit from investing is $\pi_0 < I$, which is known. The period-1 payoff from the investment is either π_L or π_H , where $\pi_H > I > \pi_L$. State H occurs with probability p_H .

Since the planning horizon is finite, we may derive the optimal strategy via backward induction. In period $t = 1$, if state H occurs, the decision maker will choose to invest ($\pi_H - I > 0$). On the contrary, if state L occurs, the decision maker chooses *not* to invest ($\pi_L - I < 0$). The optimal strategy in period $t = 1$ is therefore: invest if the state is H , do not invest otherwise. In period $t = 0$, the decision maker weighs the *expected* payoff from investing in $t = 0$ against the *expected* profit from investing in $t = 1$. She postpones investment, iff

$$p_H(\pi_H - I) - [\pi_0 - I + (p_H\pi_H + (1 - p_H)\pi_L)] > 0 \quad (3.1)$$

$$\Leftrightarrow (\pi_0 - I) + \pi_L + p_H(I - \pi_L) < 0 . \quad (3.2)$$

The first term in the first line is the continuation value of the option to invest. The second term is the stopping value from the option to invest. By collecting terms, the inequality in the second line illustrates that the lower the probability for the good state H , the more likely a decision maker is to postpone investment. Note that this effect is not completely obvious *a priori*, since a reduction in p_H affects both, the stopping and the continuation value, in the same direction.

In a more realistic setting, it seems natural to assume that p_H is not perfectly known to the decision maker. Suppose the decision maker has two possible candidates (theories) in mind $p_H \in \{.3, .7\}$ and no (or very little) objective indication which theory prevails.

Then if she adopts theory 1, investment in $t = 0$ is less likely than with theory 2. Pessimists invest later.

One could tell a similar story in terms of uncertainty aversion, however. Note that investment (or stopping) exposes the decision to uncertainty via the payoff process. Before, the decision maker is insulated against the source of uncertainty, namely the uncertainty surrounding the payoff process π_t . An uncertainty averse decision maker, who has the option not to face an uncertain payoff, shuns exposing herself to this uncertainty, relative to a person who is ambiguity neutral/seeking. Gilboa & Schmeidler (1989) provide an axiomatic foundation for such behavior in static choices, Epstein & Schneider (2003b) derive a recursive utility representation for a dynamic setting (recursive MEU). The recursive MEU representation stipulates that MEU decision maker have *multiple* prior beliefs about the probability law driving $\{\pi_t\}$, all of which they regard as equally plausible. They then behave according to the most pessimistic belief about the probability law driving $\{\pi_t\}$ they have in mind.³

From an experimental point of view, it is important to design the experiment in a way that it provides a clear separation between the two behavioral interpretations. More formally, it is the objective to identify a subject that has *multiple* prior beliefs from a subject that has a unique, but pessimistic prior belief. Without any further identification scheme, the two explanations are behaviorally equivalent.

In order to discriminate between plain pessimism and multiple priors, I use part of the experimental design applied by Ellsberg (1961). Toward that end, assume a slight variation of the above setting. There are two possible states of the world $S = \{R, B\}$ for which the decision maker has no objective prior at hand, i.e. the probability for either state is unknown (or at least very vague) *a priori*. The evolution of $\{\pi_t\}$ is then assumed to depend on the true state $s \in S$ in the following way. If the state is $s = R$ ($s = B$), then π increases from π_0 to π_H whenever a red (black) ball is drawn from an urn with 70 (30) red balls and 30 (70) black balls. Before the investment decision is to be made, let the decision maker bet on either R or B . Hence, there are two acts the decision maker may choose, r or b . Under Savage's axioms of choice, a SEU decision maker's preferences \succeq are based on beliefs, i.e. there exists a binary ordering \succeq_p such that $r \succeq_p b$ whenever r is believed to be *a priori* more likely than b . According to axiom P4 in Savage (1954), the decision maker prefers to bet on r rather than b , iff $r \succeq_p b$ (see also Epstein & Le Breton, 1993). This can be understood as bounding her prior beliefs above 0.5. Moreover, as was demonstrated by Weller (1978), a dynamically consistent SEU decision maker must be Bayesian. Under Bayesian updating, however, beliefs may not cross in the sense that if prior beliefs $\mu_0^{(i)} \geq \mu_0^{(j)}$, this relation is preserved for the posteriors $\mu_t^{(i)} \geq \mu_t^{(j)}$. In the

³What exactly is the worst-case belief may not be completely clear *ex ante* and depend on the history of $\{X_t\}$. In the simple framework used here, it is easy to identify the worst case measure *ex ante* using monotonicity and first-order stochastic dominance (see Riedel, 2009, section 4.2).

context of the simple investment decision above, if the decision maker was equipped with an objective prior on S that puts equal weight on both states, we would expect her to be at most as likely to invest as in the situation involving uncertainty and subjective priors. I demonstrate in the appendix, that this conjecture is true and that the effect prevails in a more general setting.

Under this scenario, it is then straightforward to think about an experimental design that involves a treatment group (subject to ambiguity) and a control group (subject to risk). In the first stage of the experiment, subjects are prompted to bet on the state of the world. In the second stage, they observe the realization of the profit process contingent on their choice in the first stage. Behavior can then be contrasted across groups and individuals.

4 The underlying model

This section presents the derivations of model predictions under SEU. Unlike other experimental studies on optimal stopping tasks, I derive the model predictions explicitly in discrete time. This seems to be expedient, since time is necessarily discrete in the laboratory. Moreover, I will relax assumptions made in the simple example and allow in particular for (i) infinitely many periods, (ii) risk-aversion and of course (iii) non-zero discount rates. Having derived the optimal strategy, i.e. the optimal stopping time for an SEU agent, it will become clear how to arrive at a testable hypotheses about a deviation from SEU. The notation readily generalized from that of the previous section.

4.1 The optimization problem of an SEU agent

Assume time is discrete, $t = 0, \dots$, and let the payoff process $\{\pi_t\}_{t=0}^{\infty}$ be a binomial random walk, i.e.

$$\pi_t = \begin{cases} h \pi_{t-1} & \text{with } p_H^{(s)} \\ h^{-1} \pi_{t-1} & \text{with } 1 - p_H^{(s)} \end{cases} \quad (4.1)$$

where $s \in \{R, B\}$ denotes the state of the world. Moreover, assume that $p_H^{(R)} = 1 - p_H^{(B)}$. Note that because upticks and downticks are fully reabsorbing, the state of the process in any period t is sufficiently described by the tuple (π_0, x_t) , where x_t is the number of upticks minus downticks up until period t .

The agent's objective is to maximize the value function

$$V(x_t) = \max_t \mathbb{E}_t \left[\sum_{s=t}^{\infty} \delta^{s-t} u(\pi_t h^{x_s}) \right] \quad (4.2)$$

with respect to time t . The Bellman equation for this problem is given as (see e.g. Dixit & Pindyck, 1994, chapter 4 for an accessible derivation)

$$V(x_t) = \max \{ \Omega(x_t) - I, \delta \mathbb{E} [V(x_{t+1}) | x_t] \} . \quad (4.3)$$

where

$$\Omega(x_t) = \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} u(\pi_t h^{x_s}) | x_t \right] \quad (4.4)$$

$$\mathbb{E} [V(x_{t+1}) | x_t] = p_{t+1|t} V(x_t + 1) + (1 - p_{t+1|t}) V(x_t - 1) . \quad (4.5)$$

and the $p_{t+1|t}$ are the posterior one-period ahead beliefs that the process increases by one step give some prior belief about the state of the world.

$$p_{t+1|t}(x_{t+1} = x_t + 1 | x_t) = \mu(x_t | \mu_0) p_H^{(\iota)} + (1 - \mu(x_t | \mu_0))(1 - p_H^{(\iota)}) \quad (4.6)$$

$$\mu(x_t | \mu_0) = \frac{A}{1 + A} ; A = \frac{\mu_0}{1 - \mu_0} \left(\frac{p_H^{(\iota)}}{1 - p_H^{(\iota)}} \right)^{x_t} \quad (4.7)$$

where ι is the state of world that the agent chose to bet on (let $\neg \iota$ be the opposite state) and (4.7) stipulates that the agent learns according to Bayes' rule. For later use, it is instructive to note at this point that the speed of learning depends on the difference between the probability of an uptick across states. To provide a formal argument, let ϵ denote the deviation of p_H from 0.5. Then note that in (4.7)

$$\frac{\partial \mu}{\partial \epsilon} > 0 ; \frac{\partial A}{\partial \epsilon} \propto \frac{1}{0.25 - \epsilon^2} > 0 \quad (4.8)$$

because $\epsilon \in [0, 0.5)$. If $\epsilon = 0.5$, then the benefit from learning is infinite, because a single observation is fully revealing.

Note that, if I know the function $V(\cdot)$, then I could evaluate the RHS of (4.3) for every x_t and determine the optimal strategy for each x_t (i.e. stop if the first element in the max operator is at least as large as the second element; otherwise continue). This is only true under risk-neutrality, however. Under risk-aversion, one needs to impose some form of the instantaneous utility function $u(\cdot)$. I will assume a CRRA form for $u(\cdot)$ throughout the paper. With this, I am able to exploit the Contraction Mapping Theorem (Stokey, Lucas & Prescott, 1989) to solve the problem recursively.

As I demonstrate in the appendix, the following proposition holds for the optimal strategy

Proposition 4.1 (SEU). *The optimal strategy of an agent with CRRA utility, objective function (4.2) and beliefs according to (4.7) is*

- (i) *to invest as soon as the process π_t exceeds a threshold π^* (cut-off strategy),*
- (ii) *a monotone function of her prior beliefs μ_0 , i.e. $\frac{\partial \pi^*}{\partial \mu_0} < 0$.*

The first part of the proposition is not too surprising, since the entire system may be summarized by the single state variable x_t . The second part is less obvious in the general setting. Note once more, that a reduction in μ_0 reduces $p_{t+1|t}$ for all x_t , which has a negative effect on the continuation *and* the stopping value in (4.3). The net effect is then not immediately obvious. Hence, part (ii) confirms the conjecture that the intuition from the simple example in the previous section carries over to the general framework.

Overall this shows that under SEU and everything else equal, initially more optimistic agents invest earlier than initially more pessimistic agents.

4.2 The optimization problem of an MEU agent

Assume that MEU agents face the same payoff process $\{\pi_t\}_{t=0}^\infty$, which is again a random walk as in (4.1). Conceptually, there is little difference in the derivations for the MEU agent versus the SEU agent. First, the MEU agent possesses a *set* of prior beliefs \mathcal{M}_0 on s . Following Epstein & Schneider (2003b, 2007), this set of prior beliefs is updated prior-by-prior

$$\mathcal{P}_{t+1|t} = \left\{ \mu(x_t | \mu_0) p_H^{(\iota)} + (1 - \mu(x_t | \mu_0))(1 - p_H^{(\iota)}) : \mu_0 \in \mathcal{M}_0 \right\} . \quad (4.9)$$

Second, the decision maker evaluates prospects in her objective function wrt to all measures in $\mathcal{P}_{t+1|t}$ and then takes the measure that yields minimal expected utility. How this *worst-case measure* looks is easily determined in the given case. Note that the payoff function $u(\pi_t)$ is a monotone function of the state variable x_t . Then because of this monotonicity, it should also be intuitively clear that the worst-case measure is given by

$$\mathcal{P}_{t+1|t} = \left\{ \mu(x_t | \mu_0) p_H^{(\iota)} + (1 - \mu(x_t | \mu_0))(1 - p_H^{(\iota)}) : \mu_0 = \inf \mathcal{M}_0 \right\} \quad (4.10)$$

which is a singleton set. Formally, we may argue in terms of first-order stochastic dominance to establish that this is indeed the worst-case measure. The worst-case measure is such that it assigns minimal probability to an uptick in the next period, given any history of the events x_t .

The assumptions underlying the learning process warrant some elaboration. As highlighted by Al-Najjar & Weinstein (2009), learning and updating of beliefs is a subtle issue in ambiguity models. In a general dynamic choice problem, MEU maximizer might suffer from preference reversals upon learning, i.e. they make time-inconsistent choices. More formally, if the MEU agent has an *a priori* preference relation \succeq on a set of feasible acts and $f \succ g$, his conditional preferences \succeq_E maybe such that $g \succ f$. Upon revelation of event E , the decision maker deviates from his initial plan. Epstein & Schneider (2003b) demonstrate, however, that if prior beliefs have a certain recursive structure (referred to as *rectangularity* or stability under pasting), such reversals may not occur. While this is restrictive in general, the information structure (the structure of the tree of events) in the investment experiment is such that *all* priors over the state space S are rectangular (see Riedel, 2009, appendix D for illustrative example).

The objective function of the MEU agent is then given as

$$V(x_t) = \max_t \min_{p \in \mathcal{P}} \mathbb{E}_t \left[\sum_{s=t}^{\infty} \delta^{s-t} u(\pi_0 h^{x_s}) \right] \quad (4.11)$$

$$= \max_t \mathbb{E}_{\underline{p}} \left[\sum_{s=t}^{\infty} \delta^{s-t} u(\pi_0 h^{x_s}) \right] . \quad (4.12)$$

The MEU Bellman equation is then given as

$$V(x_t) = \max \{ \Omega(x_t) - I, \delta \mathbb{E} [V(x_{t+1}) \mid x_t] \} . \quad (4.13)$$

with all expressions defined as before, but all expectations taken wrt $\underline{p}_{t+1|t}$. Consequently, the optimal strategy may be determined just as for the SEU agent, with respect to a particular measure.

4.3 Central hypothesis

As mentioned at the beginning, it is necessary to distinguish between plain pessimism induced by uncertainty and uncertainty aversion in form of multiple priors. Hence, let the experiment have two stages. During the first stage, before observing $\{\pi_t\}$, the decision maker has to decide to bet on one state of the world $s \in \{R, B\}$. If she guesses the true state correctly, $\{\pi_t\}$ will have a high uptick probability, if not it will have a low uptick probability. The idea behind this was illustrated earlier. An SEU agent in the treatment group has no *objective* prior about the likelihood that either state occurs. Given the absence of an objective prior, agents are supposed to form a *subjective* prior on S . The first-stage choice should then be dictated by their subjective prior beliefs. If $\mu_0 = Pr(s = R) \geq 0.5$ choose R , otherwise choose B . Hence, this choice can be thought as imposing a lower

bound on subjective beliefs. A corollary that can be derived from proposition 4.1 (ii), is that the bound on beliefs translates into a bound on reservation profits π^* .⁴

Corollary 4.1 (Comparative statics). *Given everything else equal, if $\mu_0^i \geq \mu_0^j$, then $\pi^{*,i} \leq \pi^{*,j}$ for agents i, j ($i \neq j$).*

This corollary gives the central hypothesis tested below. Given that we would expect subjective priors for the chosen state in the treatment group to be at least as large, we expect them to invest no later than subjects in the control group. The word 'later' here is understood in terms of the reservation profit π^* . This is illustrated in figure 1. Here, the optimal strategy for various priors of an agent with relative risk-aversion coefficient $\theta = 0.5$ is depicted.⁵

In this context, I want to highlight two points. First, it is important that subjects were told that the mechanism according to which the state of the world is selected before each round is stationary. If subjects do not know anything about how the state is determined, any behavior can be rationalized by SEU. This is because without any knowledge about the mechanism, the decision maker is not constrained to believe that the pattern is iid random. Then any (weird) belief is feasible, supporting virtually any action. Second, the probability for a state was linked to a real-world phenomenon, which subjects might perceive as something that could, in principle, be determined (with some precision). This is in contrast to the original Ellsberg experiment (and variants thereof), where there is truly no way subjects could know or find out about the probability for any composition of the ambiguous urn.

5 Implementation

This section outlines how I implemented the experimental setup in the laboratory. The experiment was conducted as a computer-based experiment. The experimental software was programmed using JScript and Python. Subjects saw two different screens in each of the 45 rounds that they played. Subjects were told that they have the option, but not the obligation, to invest in a factory, which, upon investment, produces one unit of a fictitious product every period until the end of the round.

In the laboratory, the experiment was implemented by a series of two alternating screens. The first screen prompted subjects to set a color for the coming round. The screen showed a simple radio button for each color R and B and subjects had to click a button below to confirm their choice (see fig. 4). Subjects were told that the behavior

⁴This is a direct consequence of the fact that beliefs that develop according to Bayes' rule, cannot cross for a finite set of observations. Thus, initially more optimistic agents remain more optimistic in finite time.

⁵Setting $\theta = 0.5$ was motivated by findings due to Holt & Laury (2005).

of the profits $\{\pi_t\}$ depends on whether they meet consumers' taste for the color of the product. There were two colors available, red (R) and black (B). It was then mentioned that prior to each round, they would have to set up their machines such that they produce red or black products, but not both, for the entire round. If the color they chose matched consumers' taste, the per-period profits had a 0.57 chance experiencing an uptick. If the chosen color did not match consumers taste, the probability for an uptick was $1 - 0.57 = 0.43$. The choice for the state-wise probability for an uptick versus a downtick was largely dictated in order to meet a sweet spot between the amount of uncertainty and informativeness of the realizations. First, the greater the difference between $p_H^{(t)}$ and $p_H^{(-t)}$, the greater the difference between the respective drifts of the per-period profit processes. For a negligible difference, the impact of uncertainty is supposed to be negligible as well. Hence, from an experimental point of view, a larger difference seems desirable in order to identify an effect. Second, for any value of p_H that differs appreciably from 0.5 (by more than 0.05), the drift dominates the process visually. The true state of the world is then easily discovered after only a few observations. Note that learning decreases the amount of ambiguity over time and in the limit ambiguity disappears (Marinacci, 2002, provides a formal argument). Hence, only a small difference is pertinent to maintain ambiguity for a minimum amount of time.

The second screen presented the actual investment screen. Subjects there saw the realization of a binomial random walk with parameters mentioned above. The realizations of the payoff process $\{\pi_t\}$ (measured in ECU) they observed were the (potential) per-period profits from selling the product. This process was always started at the value of 40 ECU ($\pi_0 = 40$) and the factory investment cost was fixed at 3,200 ECU. Each second consisted of two ticks (see fig. 5).

In the control group, subjects were informed that the probability for each color to be the correct one was 0.5.⁶ In the treatment group, instructions were the same as in the control group, apart from the information concerning the prior probability for either state of the world. Subjects in the treatment group, were told that the probability for red to be consumers' taste was equal to the average relative amount of rainy days per year in Jakarta (Capital of the Republic of Indonesia) between 1971 and 2000. Unlike in related experimental papers (Oprea et al., 2009; Della Seta et al., 2012) subjects were not

⁶Note that this is the most conservative choice in terms of risk aversion. It is easy to demonstrate that risk aversion has the same effect as pessimism, i.e. higher risk aversion increases π^* . For a given stepsize h , the posterior variance is only a function of the binomial variance of the increments of $\{\pi_t\}$

$$\sigma^2 \propto p_{t+1|t}(1 - p_{t+1|t}) \quad (5.1)$$

This expression has its maximum at 0.5 and decreases symmetrically from there. For any given history x_t , it is closest to 0.5, if $\mu_0 = 0.5$. If now the posterior variance σ^2 of the process $\{\pi_t\}$ enters $u(\cdot, \sigma^2)$, this makes it more likely that the agent postpones investment. Hence, it becomes less likely to observe a mark-up on reservation profits π^* in the treatment group relative to the control group.

constrained to seize the investment at a value strictly larger than the initial value or the investment costs. A round ended randomly, with a given and constant probability of 0.7% (i.e. on average $t = 143$ ticks).⁷

The 45 random walks used in the experiment were the same for every subject. Based on an individual login (printed on the instructions), the set of 45 random walks was stratified over 45 rounds.⁸ The experiment was designed so that for each subject in the treatment group, there was one subject in the control group that saw the same sequence over 45 rounds (contingent on choosing the same color in a given round). This way, a potential 'round effect' is supposed to be mitigated. Otherwise subjects might be framed by particularly short/long realizations of the process in the first few rounds.

The experiment was conducted at the Technical University Berlin (TU) and the WZB Berlin. The preliminary dataset presented here, was obtained from three laboratory sessions with randomly selected students from the ORSEE pool of the TU and WZB.⁹ For each session 22 students from various fields of study participated in 45 rounds of the experiment for pay. Participants were randomly selected into either the treatment or the control group (as above: treatment=ambiguity, control=risk). Consequently, each group currently has 33 subjects with 45 observations each (potentially right censored). The average duration of the experiment was 74 minutes, and the mean earnings for subjects was 15.98 Euros (median=14 Euros), where the minimum and the maximum payment were 5.00 Euros and 39.00 Euros respectively.

6 Experimental results

This section outlines the results obtained from the first three sessions, all currently available data. The analysis is conducted in view of the central hypothesis derived in section 4.3. It employs statistical methods from the field of Survival or Duration Analysis to surmount two key data defections, i.e. the effect of censoring, and unobserved heterogeneity across subjects.

6.1 The effect of ambiguity

Due to the termination hazard, 37% of the observations in the sample are right censored. In these cases, the process ceased before a subject decided to seize the investment. Ignoring this fact introduces a censoring bias into standard estimators, while dropping censored

⁷A translated version of the instructions is available from the author upon request.

⁸The login name consists of an initial, "A" or "B", and a two-digit number. The letter prefix indicates which treatment a subject belongs to, while the two digit number was used as the seed for a pseudo-random number generator that drew the sequence of series shown. In both groups, subjects were then numbered in an increasing order.

⁹See: <https://experimente.wzb.eu/>

observations results in a truncated sample and leads to a truncation bias. In order to overcome this issue, this section makes use of tools from Survival Analysis, which provides unbiased estimates in the presence of censoring.

Evidence from Kaplan-Meier estimation

Following Oprea et al. (2009); Della Seta et al. (2012), results are first analyzed by group using a non-parametric Kaplan-Meier (Kaplan & Meier, 1958) estimator. This estimator focuses on the distribution of the reservation profit. It estimates the survival function, which corresponds to the probability *not* to invest at a given value of the profit process. If N_π is the number of subjects who did not invest at a value of π *excluding* those for which the process terminated at that value, and Y_π is the number of subjects who invest at a given value of π , then the Kaplan-Meier estimator of the Survival Function is defined as

$$\hat{S}(\pi) = \prod_{\pi=0}^{\pi_{max}} \left(1 - \frac{Y(\pi)}{N(\pi)} \right) . \quad (6.1)$$

The idea behind the Kaplan-Meier estimate is to provide a standard empirical distribution function of reservation profits π^* , taking into account that at various instances, subjects drop out of the set of subjects that still have the opportunity to invest. In the absence of censoring, (6.1) coincides with the empirical distribution function of π . On the one hand, the Kaplan-Meier estimate provides a way to determine the direction of the effect of ambiguity, but is less suitable for gauging the magnitude of the effect. On the other hand, this procedure is truly non-parametric, hence quite robust against misspecification (e.g. see Therneau & Grambsch, 2000, chapter 2 for an in-depth discussion).

Figure 2 shows the estimated survival function by group. As shown, the survival functions for both groups separate in a direction that contradicts the SEU model prediction. Instead, subjects in the treatment group tend to react to uncertainty in a way that is predicted by the MEU model. For a given value of the per-period profit process $\{\pi_t\}$, subjects in the treatment group have a lower probability to seize the investment. Note that under the hypothesis that subjects in the treatment group are SEU maximizer, we would expect the opposite.

One may perform a statistical test for equality of the two cdfs, by means of a log-rank test (Harrington & Fleming, 1982) with the Nullhypothesis

$$H_0 : S_1(t) = S_2(t) . \quad (6.2)$$

The associated statistic is $\chi^2(1)$ distributed. The value of the statistic is 45.76, with an associated p-value which is virtually zero.

The impact of uncertainty may be further quantified, by considering the average reservation profit within each group. Table 1 displays the estimates for the mean and the

	Mean	SE	Median	95% int.
Risk group	64.28	1.02	41.2	[41.2, 42.44]
Ambiguity group	72.57	1.1	45.02	[45.02, 46.37]
no. of investments risk group			1017	
no. of investments ambiguity group			859	

Table 1: Mean and median reservation profit by group (N=2970, upper limit 122.99134).

median reservation profit by group. The estimate for the mean duration suffers from a bias that stems from the fact that the survival function does not become zero over the feasible state space. Consequently, the integral has to be cut off at the highest censoring value. Comparing median and mean estimates, the bias seems to be very pronounced in the given case. The median, however, may still be estimated consistently from the data. In terms of the median reservation profit, the magnitude of the difference between reservation profits is less striking and is roughly equal to a 9.3% increase in the reservation wage. Considering the 95% confidence intervals around the medians shows, however, that the difference is significant.

Evidence from mixed proportional hazard models

In an attempt to gauge the size of the effect, regression-based methods provide a more adequate tool. In the field of survival analysis, there are several models available, which mainly differ in the amount of parametric rigor they impose on the functional form of the so-called *hazard function* $\lambda(\pi)$. The hazard function and the survival function are related by the equation

$$S(\pi) = \exp \left[- \int_0^\pi \lambda(p) dp \right] . \quad (6.3)$$

Hence, the hazard function is the instantaneous probability to invest at a given value of π . Proportional hazard models assume that the hazard function for individual i comprises the *baseline hazard function* and the *risk score*. It furthermore assumes that baseline hazard and risk score are related in a proportional way

$$\lambda(\pi) = \lambda_0(\pi) \exp [X_i \beta] \quad (6.4)$$

where X_i is the i -th row of the $(n \times k)$ -matrix of covariates and $\exp [X_i \beta]$ is the risk score. In this section, I employ models that treat the baseline hazard $\lambda_0(\pi)$ non-parametrically (see Cox, 1972), going by the name Cox proportional hazard models. In such models, the baseline function $\lambda_0(\pi)$ is allowed to have any shape. In the present case, X_i comprises

	coefficient	exp(coeff.)	SE	z-stat.	Pr(> z)
treatment	-0.44	0.65	0.23	-1.93	0.054
Rand. effect		Std. dev.			Variance
		0.76			0.58
Null model likelihood			-14084.86		
Integrated partial likelihood			-13780.17		
LR test for random effects			-558.06	p: 0.00	

Table 2: Results from mixed effects Cox model (N=2970, no. of investments = 1876). Efron approximation for ties.

only a dummy variable for being a member of the treatment group. It is well-known, however, that unobserved heterogeneity among individuals in a study leads to inconsistent estimates for the treatment fixed effect β in (6.4) (e.g. Chamberlain, 1984, 1985; Lancaster, 1985; Aalen, 1988). Because I do not have further controls to include in X_i , a great number of individual characteristics are unobserved which potentially influence decisions. Consequently, the basic Cox model is extended by incorporating an individual-specific random effect θ (often called frailty effect), to absorb the unobserved heterogeneity into the risk score

$$\lambda_i(\pi) = \lambda_0(\pi) \exp [X_i\beta + Z_i\theta] . \quad (6.5)$$

It is assumed here that random effects are normally distributed

$$\theta \sim N(0, \sigma^2 I_n) . \quad (6.6)$$

The model (6.5)-(6.6) may be estimated using penalized regression methods (see Hastie & Tibshirani, 1990; Therneau, 2003).

Results from the Cox model with Gaussian frailties are given in table 2. The results may be easily interpreted in terms of the relative risk score, which is the probability for a subject to invest at a given value π , relative to a subject in the control group. Note that conditional on θ the *relative* hazard for a member of group i is given as

$$\frac{\lambda_{i=1}(\pi)}{\lambda_{i=2}(\pi)} = \frac{\lambda_0(\pi) \exp [X_1\beta]}{\lambda_0(\pi) \exp [X_0\beta]} = \exp [(X_1 - X_0)\beta] = \exp [\beta] . \quad (6.7)$$

Table 2 then reveals that the conditional relative risk score is 0.65, i.e. given the individual frailty term, ambiguity reduces the probability to invest at a given level by around 35%. However, the effect is borderline significant ($p = 0.054$). The results also show that the variance of the random effect is significantly larger than zero, as confirmed by an LR

test.¹⁰ The estimated variance of θ is substantial. Deviating one standard deviation from the mean of θ (which is zero by assumption), translates into 2.14 times the probability to invest at a given price. Hence, individual heterogeneity is found to be highly significantly different from zero. The magnitude of the dummy for being in the treatment group makes the treatment effect appear quite important as well. Concerning its statistical significance, however, reveals that it is only marginally significant at a level of 5.4%.

The overall results illustrate the importance to take into consideration the presence of individual heterogeneity that is due to unobserved factors.

7 Conclusion

This paper considers the impact of ambiguity or Knightian uncertainty on investment decisions that are irreversible. Toward that end, the paper presented a theoretical model based on from which model predictions under expected utility (Dixit & Pindyck, 1994) and multiple-prior expected utility (Nishimura & Ozaki, 2007; Riedel, 2009) were derived. In a computer-based laboratory experiment, subjects then had to repeatedly decide whether or not they want to seize an investment opportunity. Upon investment, they earned the sum of future values of a stochastic payoff process.

The underlying theoretical model predicts that ambiguity averse decision makers react to ambiguity by postponing the investment, relative to a situation where there is risk. The main finding of this paper is that data obtained from an experiment with a treatment group (subject to ambiguity) and a control group (subject to risk) confirms this prediction. A non-parametric analysis of observed reservation profits revealed that subjects in the treatment group have a 9% higher median reservation profit than subjects in the control group. A regression-based analysis across groups that controls for unobserved heterogeneity, revealed that subjects in the treatment group are 0.65 times as likely as subjects from the control group to seize the investment at a given level of the payoff process.

These results indicate that uncertainty indeed plays a role for an individual's decision to pledge a safe payoff in favor of an uncertain payoff. Subjects in the laboratory ask for a higher reservation profit to forego a safe endowment for an uncertain investment. Even though, interpreting experimental results outside of the laboratory context must be handled with care, one may still think how this is important in every-day life. For instance, on top of a regular volatility premium due to risk aversion, experimental results indicate the existence of an ambiguity premium. Consequently, investors may shun taking

¹⁰A Cox model without frailties was also estimated. In such a model, the treatment effect was significant at a 2.5% level and the relative risk score was 0.75. The proportionality assumption (6.4), however, was strongly rejected, as indicated by the scaled Schoenfeld residuals.

positions as market conditions become more uncertain. While this is often interpreted as a decrease in risk-appetite, this may also stem from heightened uncertainty without any increase in fundamental volatility. This result is even more striking because the amount of ambiguity and complexity in the experiment may be considered as fairly small compared to real-world situations. Consequently, it seems possible that the estimated effect in this paper understates potential ambiguity premium in real-world investment decisions.

Appendices

Technical Appendix

This appendix discusses the details of the Value Function Iteration algorithm mentioned in section 4.1. Start with the Bellman equation (4.3) for the optimization problem

$$V(x_t) = \max \{ \Omega(x_t) - I, \delta \mathbb{E} [V(x_{t+1}) \mid x_t] \} . \quad (7.1)$$

(7.1) is equivalent to

$$V(x_t) = \max_{v_t} u(\pi_0 h^{x_t}, v_t) + \delta [p_{t+1|t} V(x_t + 1) + (1 - p_{t+1|t}) V(x_t - 1)] \quad (7.2)$$

where v_t is the binary control variable, $v_t \in \{0, 1\}$ and 1 = 'invest' and 0 = 'postpone'. The latter form will be more convenient to use below. I seek a function $V(\cdot)$ that solves (7.1) for every x_t . With such a function, the RHS can be evaluated on the state space and the set of states for which stopping is optimal (the active set) may be found. Following Stokey et al. (1989), equation (7.1) may be viewed as a functional equation on the space (V, d_∞) , i.e. the space of bounded functions equipped with the sup-metric. Since time is supposed to be infinite, one may view the state space as independent of time (it merely denotes the current value of the per-period profit process). Moreover, understand the RHS of (7.2) as an operator, i.e. a mapping $\Psi : V \rightarrow V$

$$V(x) = \Psi(V(x')) \quad (7.3)$$

where in the binary case here x' is either $x - 1$ or $x + 1$. Using Blackwell's sufficient conditions (see e.g. Ljungqvist & Sargent, 2004, appendix A.1), one may prove that in fact Ψ , as defined by (7.1) is a contraction with modulus δ . First, for every two functions V and W such that $V > W \forall x$ it holds that

$$\Psi(V) = \max_{v_t} u(\pi_0 h^{x_t}, v_t) + \delta [p_{t+1|t} V(x_t + 1) + (1 - p_{t+1|t}) V(x_t - 1)] \quad (7.4)$$

$$\geq \max_{v_t} u(\pi_0 h^{x_t}, v_t) + \delta [p_{t+1|t} W(x_t + 1, v_{t+1}) + (1 - p_{t+1|t}) W(x_t - 1, v_{t+1})] \quad (7.5)$$

$$= \Psi(W) \quad (7.6)$$

and for any positive constant k

$$\Psi(V + k) = \Psi [u(\pi_0 h^{x_t}) + \delta [p_{t+1|t} (V(x_t + 1) + k) + (1 - p_{t+1|t}) (V(x_t - 1) + k)]] \quad (7.7)$$

$$= \Psi(V) + \delta k . \quad (7.8)$$

Therefore $V(\cdot)$ is shown to be monotone and to be discounting. Then by the contraction mapping theorem, there exists a unique $v^* \in V$, such that

$$\Psi(v^*) = v^* . \quad (7.9)$$

This element v^* may be found recursively via

$$v_0 \in V \quad (7.10)$$

$$v_{i+1} = \Psi(v_i) \quad (7.11)$$

and the sequence $\{v_i\}$ converges uniformly to v^* . This recursive algorithm is called Value Function Iteration.

To be able to employ the algorithm here, we have to determine the stopping value in equation (7.1) for each tuple (x, x') . This is done here under the assumption of a CRRA instantaneous utility function $u(\pi) = \pi^\theta$

$$V(x_t) = \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} u(\pi_0 h^{x_t} h^{x_s}) \right] \quad (7.12)$$

$$= \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} (\pi_t h^{x_s})^\theta \right] \quad (7.13)$$

$$= h^\theta \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} u(\pi_t h^{x_{s-1}})^\theta \right] \quad (7.14)$$

$$= h^\theta V(x_t - 1) . \quad (7.15)$$

Consequently, the solution $V(\cdot)$ is of the form

$$V(x_t) = \beta h^{\theta x_t} . \quad (7.16)$$

I may use

$$V(x_t) = u(\pi_0 h^{x_t}) + \delta [p_{t+1|t} V(x_t + 1) + (1 - p_{t+1|t}) V(x_t - 1)] \quad (7.17)$$

as an additional condition to determine the unknown coefficient β .

$$\beta h^{\theta x_t} = (\pi_0 h^{x_t})^\theta + \delta [p_{t+1|t} (\beta h^{(x_t+1)\theta}) + (1 - p_{t+1|t}) (\beta h^{(x_t-1)\theta})] \quad (7.18)$$

$$\beta = \pi_0^\theta + \beta \delta [p_{t+1|t} (h^\theta) + (1 - p_{t+1|t}) (h^{-\theta})] . \quad (7.19)$$

This implies that

$$\beta = \frac{A}{1-B} ; A = \pi_0^\theta ; B = \delta [p_{t+1|t}h^\theta + (1-p_{t+1|t})h^{-\theta}] . \quad (7.20)$$

Equations (7.16) and (7.20) then give a closed-form solution for the stopping value at each point x_t .

The algorithm was initialized with the constant function $V(x) = I = 3,200$. Convergence was achieved after several hundred iterations (~ 800).

Figures

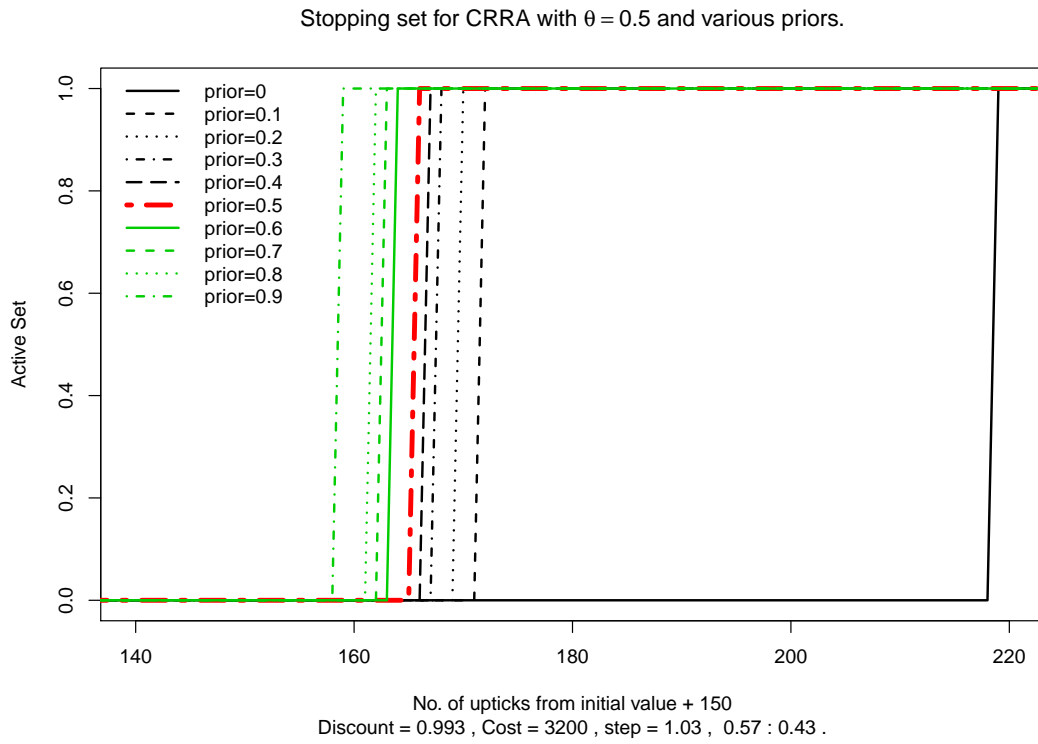


Figure 1: Active sets for an agent with various priors. The continuation region is where the step-function is zero. The stopping region is where the step-function is equal to one.

Survival function for both groups

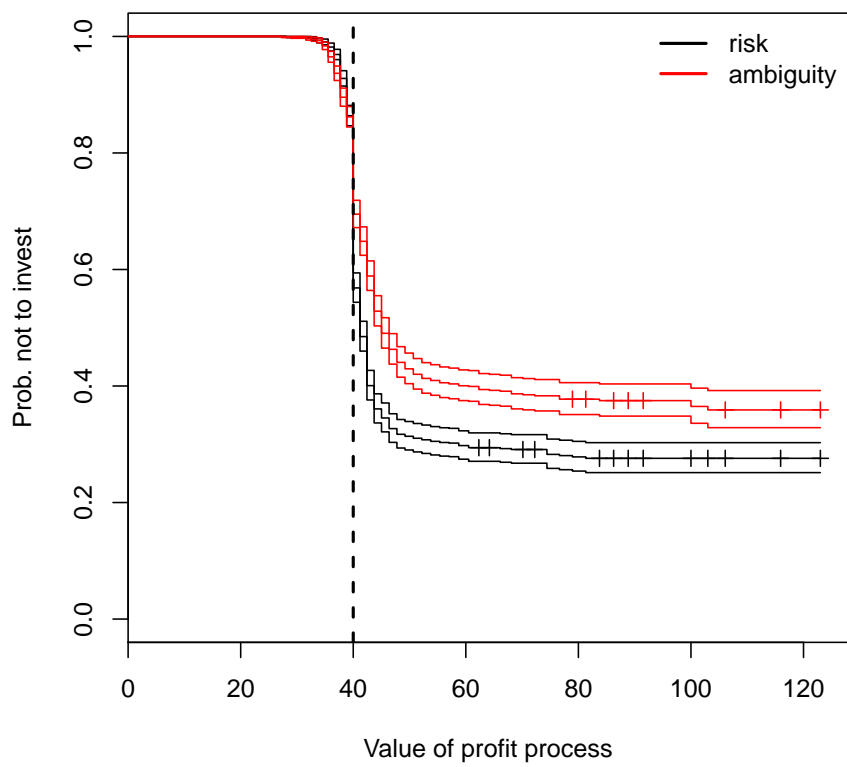


Figure 2: Estimate of the survival function by group. Standard errors are not clustered by subject. The dashed vertical line indicates the initial value of the payoff process.

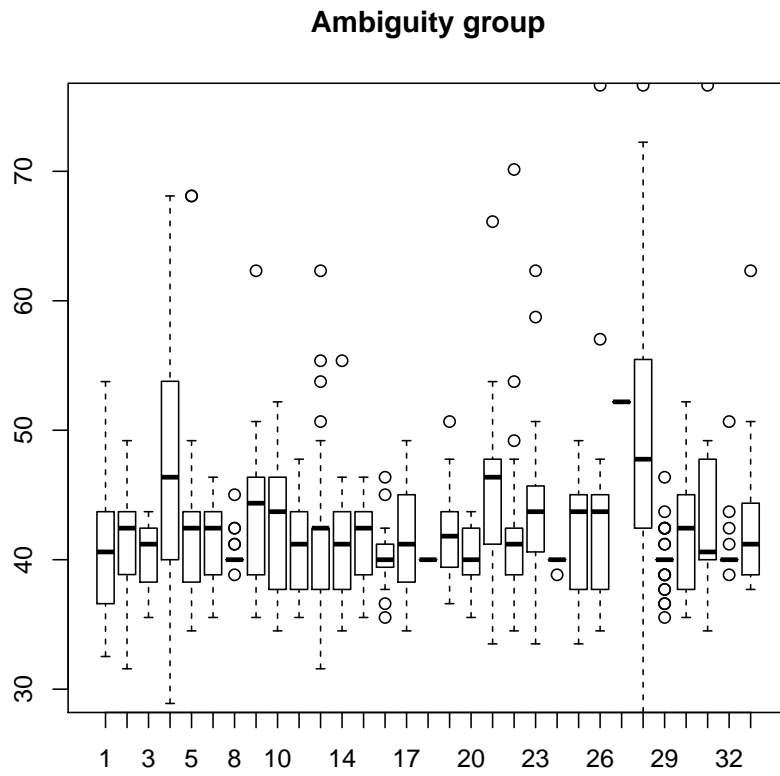
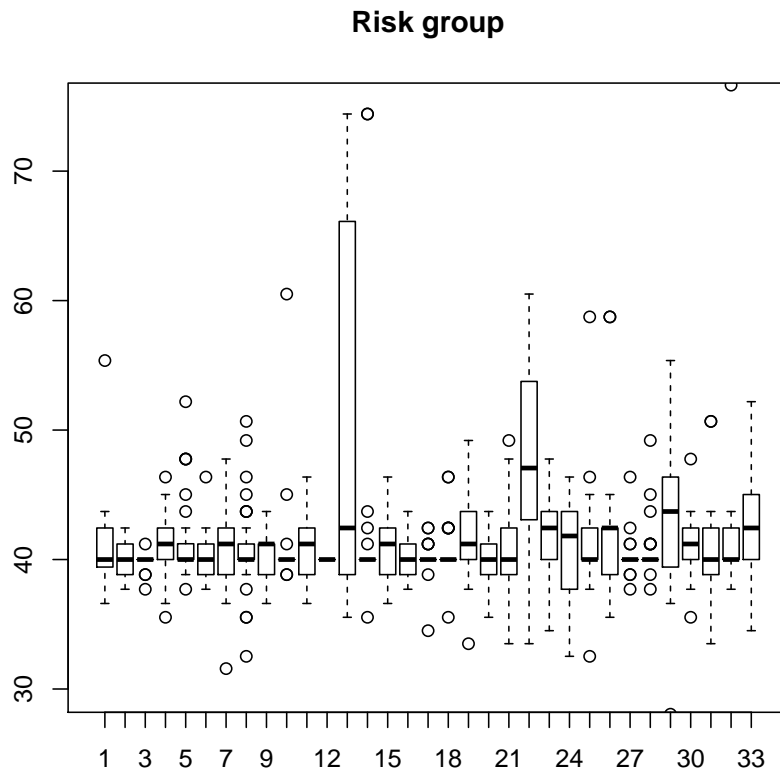


Figure 3: Reservation profits over 45 rounds by subject.

Bitte wählen Sie auf welche Farbe die Produktion in der nächsten Runde eingestellt werden soll:

- Auf Rot
- Auf Schwarz

Figure 4: Screen for color choice (in German).

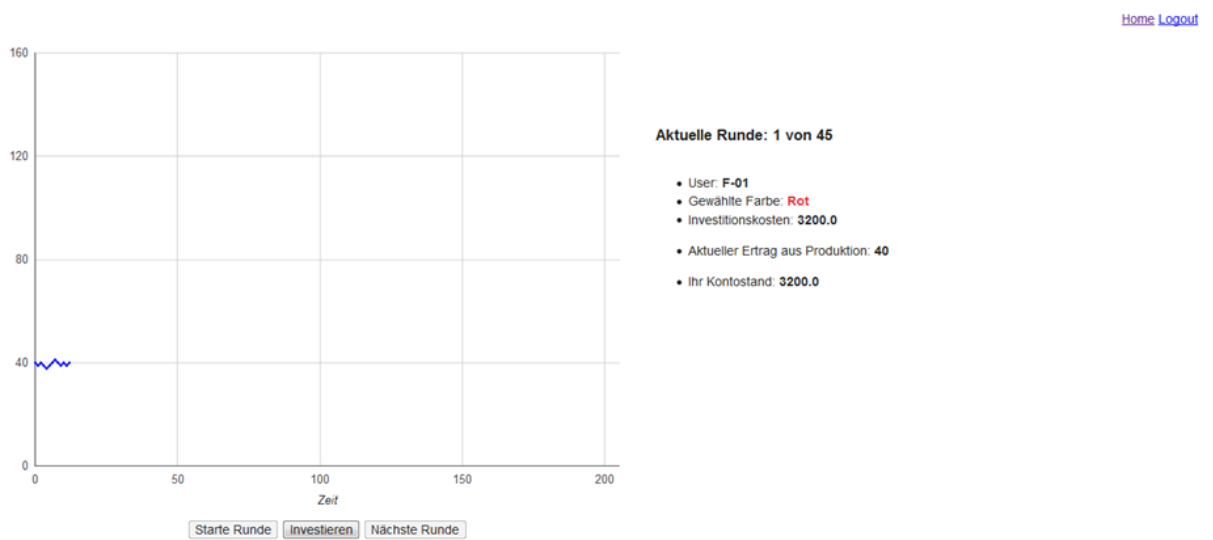


Figure 5: Investment screen (in German).

References

- Aalen, O. (1988). Heterogeneity in survival analysis. *Statistics in Medicine*, 7(11), 1121–1137.
- Al-Najjar, N. & Weinstein, J. (2009). The ambiguity aversion literature: a critical assessment. *Economics and Philosophy*, 25(3), 249–284.
- Arora, S., Barak, B., Brunnermeier, M., & Ge, R. (2011). Computational complexity and information asymmetry in financial products. *Communications of the ACM*, 54(5), 101–107.
- Artzner, P., Delbaen, F., Eber, J., & Heath, D. (1999). Coherent Measures of Risk. *Mathematical finance*, 9(3), 203–228.
- Asano, T., Okudaira, H., & Sasaki, M. (2011). An experimental test of a search model under knightian uncertainty. *Discussion Papers in Economics and Business*.
- Camerer, C. & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4), 325–370.
- Chamberlain, G. (1984). Panel data. In Z. Griliches & M. Intriligator (Eds.), *Handbook of Econometrics Vol. 2* (pp. 1248–1313). Amsterdam: North-Holland.
- Chamberlain, G. (1985). Heterogeneity, duration dependence and omitted variable bias. In J. Heckman & B. Singer (Eds.), *Longitudinal Analysis of Labor Market Data*. Cambridge: Cambridge University Press.
- Chow, Y., Robbins, H., & Siegmund, D. (1971). *Great expectations: The theory of optimal stopping*. Houghton Mifflin.
- Cox, D. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 187–220.
- Cox, J. & Oaxaca, R. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, 2(3), 301–329.
- Cox, J. & Oaxaca, R. (2000). Good news and bad news: Search from unknown wage offer distributions. *Experimental Economics*, 2(3), 197–225.
- Della Seta, M., Gryglewicz, S., & Kort, P. (2012). Willingness to Wait Under Risk and Ambiguity: Theory and Experiment.
- Dixit, A. & Pindyck, R. (1994). *Investment under Uncertainty*. Princeton University Press.
- Dominiak, A., Dürsch, P., & Lefort, J. (2009). A dynamic Ellsberg urn experiment. University of Heidelberg Working Paper.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 643–669.

- Epstein, L. & Le Breton, M. (1993). Dynamically consistent beliefs must be bayesian. *Journal of Economic theory*, 61, 1–1.
- Epstein, L. & Schneider, M. (2003a). IID: independently and indistinguishably distributed. *Journal of Economic Theory*, 113(1), 32–50.
- Epstein, L. & Schneider, M. (2003b). Recursive multiple-priors. *Journal of Economic Theory*, 113(1), 1–31.
- Epstein, L. & Schneider, M. (2007). Learning under ambiguity. *The Review of Economic Studies*, 74(4), 1275–1303.
- Epstein, L. & Schneider, M. (2010). Ambiguity and Asset Markets. *Annual Review of Financial Economics*, 2(1), 315–346.
- Etner, J., Jeleva, M., & Tallon, J. (2009). Decision theory under uncertainty. Documents de Travail du Centre d'Économie de la Sorbonne.
- Föllmer, H. & Schied, A. (2004). *Stochastic finance* (2nd ed.), volume 27 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, extended edition.
- Gilboa, I., Postlewaite, A., & Schmeidler, D. (2008). Probability and uncertainty in economic modeling. *The Journal of Economic Perspectives*, 22(3), 173–188.
- Gilboa, I. & Schmeidler, D. (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18(2), 141–154.
- Guidolin, M. & Rinaldi, F. (2010). Ambiguity in asset pricing and portfolio choice: a review of the literature. *Federal Reserve Bank of St. Louis Working Paper Series*.
- Halevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica*, 75(2), 503–536.
- Hansen, L. & Sargent, T. (2001). Robust control and model uncertainty. *The American Economic Review*, 91(2), 60–66.
- Harrington, D. & Fleming, T. (1982). A class of rank test procedures for censored survival data. *Biometrika*, 69(3), 553–566.
- Hastie, T. & Tibshirani, R. (1990). *Generalized Additive Models*. Chapman & Hall/CRC.
- Holt, C. & Laury, S. (2005). Risk aversion and incentive effects: New data without order effects. *The American Economic Review*, 95(3), 402–404.
- Kahan, J., Rapoport, A., & Jones, L. (1967). Decision making in a sequential search task. *Attention, Perception, & Psychophysics*, 2(8), 374–376.
- Kaplan, E. & Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 457–481.

- Knight, F. H. (1921). Risk, Uncertainty and Profit, 1921. *Boston and New York*.
- Lancaster, T. (1985). Generalised residuals and heterogeneous duration models: With applications to the Weibull model. *Journal of Econometrics*, 28(1), 155–169.
- Ljungqvist, L. & Sargent, T. (2004). *Recursive macroeconomic theory* (2nd ed.). Cambridge, Massachusetts: The MIT Press.
- Marinacci, M. (2002). Learning from ambiguous urns. *Statistical Papers*, 43(1), 143–151.
- McDonald, R. & Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4), 707.
- Miao, J. & Wang, N. (2011). Risk, uncertainty, and option exercise. *Journal of Economic Dynamics and Control*, 35(4), 442–461.
- Mukerji, S. & Tallon, J. (2004). An overview of economic applications of david schmeidler's models of decision making under uncertainty. *Uncertainty in economic theory: essays in honor of David Schmeidler's 65th birthday*, 63, 283.
- Nishimura, G. & Ozaki, H. (2007). Irreversible investment and Knightian uncertainty. *Journal of Economic Theory*, 136(1), 668–694.
- Nishimura, K. & Ozaki, H. (2004). Search and knightian uncertainty. *Journal of Economic Theory*, 119(2), 299–333.
- Oprea, R., Friedman, D., & Anderson, S. (2009). Learning to wait: A laboratory investigation. *Review of Economic Studies*, 76(3), 1103–1124.
- Rapoport, A. & Tversky, A. (1966). Cost and accessibility of offers as determinants of optional stopping. *Psychonomic Science*.
- Rapoport, A. & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5(2), 105–120.
- Riedel, F. (2004). Dynamic coherent risk measures. *Stochastic processes and their applications*, 112(2), 185–200.
- Riedel, F. (2009). Optimal stopping with multiple priors. *Econometrica*, 77(3), 857–908.
- Riedel, F. (2010). Optimal stopping under ambiguity in continuous time.
- Savage, L. (1954). *The Foundations of Statistics* (2nd ed.). New York: John Wiley and Sons.
- Schotter, A. & Braunstein, Y. (1981). Economic search: an experimental study. *Economic Inquiry*, 19(1), 1–25.

- Seale, D. & Rapoport, A. (1997). Sequential Decision Making with Relative Ranks: An Experimental Investigation of the. *Organizational Behavior and Human Decision Processes*, 69(3), 221–236.
- Snell, J. (1952). Applications of martingale system theorems. *Trans. Amer. Math. Soc.*, 73(293-312).
- Stokey, N. L., Lucas, R. E., & Prescott, E. C. (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Therneau, T. (2003). Penalized survival models and frailty. *Journal of Computational and Graphical Statistics*, 12(1), 156–175.
- Therneau, T. & Grambsch, P. (2000). *Modeling survival data: Extending the Cox Model*. Springer.
- Weller, P. (1978). Consistent intertemporal decision-making under uncertainty. *The Review of Economic Studies*, 45(2), 263–266.