MRIC | WORKING PAPERS

www.mric.lmu.de/research/wp

Learning and Risk Management Decisions

Dominik Lohmaier

Munich Risk and Insurance Center

Working Paper 26

July 1, 2015

Learning and Risk Management Decisions
Dominik Lohmaier
MRIC Working Paper No. 26
July 2015

ABSTRACT

Due to the lack of descriptive information about the effectiveness of risk management activities, decision-makers often have to rely on (their own) prior experience with these investments. Thus, we propose a novel, feedback-based approach to examine risk management decisions. We simulate individuals' decisions over 50 time periods and analyze how distributional properties of different risk management instruments influence subjects' propensity to invest in self-insurance or self-protection. Our results show that subjects act more risk averse over time because self-insurance take-up rates increase when learners gain more experience. Individuals' risk management decisions are price-sensitive but we find only limited evidence for a complementary relationship between both risk management options. Our findings provide an alternative explanation for the low demand for risk management investments against low probability risks and can be used to predict developments in new insurance markets with inexperienced policyholders.

Keywords: Learning from experience · Risk Management · Low probability risks

JEL Classification: D81 · D83

Dominik Lohmaier
Munich Risk and Insurance Center
Munich School of Management
Ludwig-Maximilians-Universität in Munich
Schackstraße 4/III
80539 Munich
Germany
lohmaier@bwl.lmu.de

I am thankful for helpful comments from participants at MRIC Research Seminar 2015 and the German Insurance Science Association Meeting 2015. I am indebted to Johannes Jaspersen, Stefan Neuß and Richard Peter for valuable discussions. All remaining errors are mine.

Motivation

Individuals and companies make a wide variety of risk management decisions on a regular basis. Expected utility theory usually assumes that decision-makers have comprehensive descriptions of loss distributions and of how these distributions are altered by risk management decisions (e.g., Ehrlich & Becker, 1972; Dionne & Eeckhoudt, 1985; Briys & Schlesinger, 1990; see Courbage et al., 2013 for a comprehensive overview).

However, for an individual, it is often difficult to assess to what extent his loss exposure is affected by risk management decisions, such as the installation of smoke detectors or fireproof doors. An intuitive way to approach such decision problems is to rely on prior experience. Interestingly, experimental evidence shows that subjects act very differently when they make decisions based on their experience from when they make decisions based on descriptions. The description-experience gap is even more pronounced when people face low-probability risks (Hertwig et al., 2004; Erev & Barron, 2005). For instance, Hertwig et al. (2004) show that people behave as if they underweight rare events when they make decisions based on prior experiences. The authors argue that decision-makers rely on small information samples and overweight recent events. In contrast, subjects tend to overweight small probabilities when they receive a description of the same decision problem (e.g, Kahneman & Tversky, 1979; Tversky & Kahneman, 1992; Barron & Erev, 2003). The difference between description- and experience-based decisions is of particular importance for risk management investments against low-probability risks. However, prior literature usually assumes that decision-makers use descriptive information for their risk management decisions. Hence, we propose a novel, feedback-based approach to analyze risk management decisions that aims to complement the existing risk management literature.

We develop an experiential learning model (Bush & Mosteller, 1955; Luce, 1959; March, 1996) to examine how decision-makers' propensity to invest in self-protection and self-insurance varies depending on the probability of loss and on prior experience. Our model does not impose any preferences for a certain option explicitly. Subjects are initially ignorant with respect to the underlying loss distribution and base subsequent risk management decisions on their experience. We simulate individuals' decisions over 50 time periods in order to analyze how the distributional properties of different risk management options influence subjects' probability to invest in self-insurance or in self-protection.

In line with prior research on learning and risk taking (e.g. March, 1996; Denrell & March, 2001) our simulation-based results show that subjects act more risk averse over time because self-insurance take-up rates increase when learners gain more experience. Moreover, experiential learners increase self-insurance take-up rates when they face risks with higher probabilities of loss. In addition, we are able to capture the well-known empirical phenomenon of low demand for insurance against low probability, high consequence risks in our simulative study. Interestingly, the impact of learning on self-protection take-up rates is more ambiguous. We observe the highest take-up rates for probabilities of loss of about 0.04 and the probability of investing in self-protection increases over time only for low to medium (0.005-0.05) levels of probability of loss.

By varying the cost of both risk management options, we are able to show that learners' risk management decisions are price-sensitive for most probabilities of loss. If the probabilities of loss are too small, varying the price of a risk management instrument does not affect decision-makers' choice at all because the impact on the expected value of the respective option is too small. Our evidence for cross-price effects (substitutive or complementary relationships between self-insurance and self-protection) is less clear. Learners behave as if they regard both options as complements when we vary the cost of self-protection for probabilities of loss between 0.05 and 0.2. However, we do not observe a similar relation for lower levels of probability of loss. At low probabilities of loss, we detect a weak substitutive relationship when altering the cost of self-insurance.

Our paper complements the literature with a novel approach for analyzing risk management decisions over time which requires fewer assumptions about preferences and information sets of the decision-makers than other economic models like Expected Utility or Prospect Theory. We also contribute to the literature of learning and risk taking by extending the analysis of risk taking decisions beyond the simple tradeoff between risky and safe options to the wide variety of different risk management decisions. Finally, we offer an alternative explanation for the low demand for risk management options against low probability, high consequence risks – the inability to learn from rare events. Because our model is fairly general and relies only on a few basic assumptions, our findings can be used to derive recommendations for a broad range of applications: for instance, establishing new insurance markets for inexperienced policyholders, e.g. microinsurance or developing new insurance products for emerging risks, such as cyber risks or nanotechnology.

The paper proceeds as follows. Section 2 provides an overview of the existing literature on learning and risk management and shows that this paper links two, former fairly unrelated strands of literature. We outline our model of reinforcement learning in section 3 and describe the impact of risk management on the properties of the loss distribution in section 4. Section 5 outlines our different simulation runs. We present our results in section 6 and show a sensitivity analysis in the subsequent section. We conclude with a discussion of the implications of our findings.

Learning and Risk Management

March (1988) and March and Shapira (1992) argue that individuals' and organizations' risk preferences are not constant over time but could depend on the context of the decision problem and on prior experience. Thus, the evaluation of an option's payoff could be affected by a decision-maker's goals, often coined aspiration levels. As a consequence, risk preferences may vary depending on the outcomes of a decision-maker's actions relative to a reference point. Importantly, aspiration levels are not fixed but can be adjusted depending on observed outcomes.

March (1996) analyzes how individuals' risk-taking behavior is altered by feedback about the outcome of an option relative to their aspiration levels. Decision-makers in his simulation model can choose between one riskless option offering a certain payoff and a risky option with a binomial distribution that provides the same expected value. March's results show that experiential learning leads to a rejection of the risky alternative in the gain domain but favors risky options in the loss domain, at least in the short run.[1] Thus, sequential sampling may result in an underestimation of the actual value of risky options, particularly in the gain domain. March concludes that this risk-averse behavior may be the result of learning and does not necessarily reflect utility functions or human traits.

Denrell and March (2001) extend prior simulation-based studies by examining how learning affects decisions between a risk-free option and a risky option for which payoffs are determined by a symmetric and continuous distribution. Negative outcomes in early stages could lead to the rejection of risky options, although they might provide a higher expected value. Thus, their results indicate that experiential learners act as if they were risk averse.

Denrell (2007) complements prior findings by developing a formal model on sequential sampling and risk taking. Interestingly, he shows that decision-makers exhibit risk-averse behavior only if the payoff distribution of the risky alternative is symmetric. Adaptive sampling may even generate risk-seeking behavior when the distribution of the uncertain alternative is negatively skewed. In contrast to March (1996), Denrell shows that the propensity to take risks does not depend on the payoff domain.

Jaspersen and Peter (2014) elaborate on the role of skewness preferences in experiential learning models. Because it is not possible to disentangle the effect of variance and skewness on subjects' propensity to take risks with March's lotteries, the authors develop a simulative model that enables the analysis of skewness preferences while holding the variance of the outcome distribution constant. They show that experiential learning implies skewness-averse behavior. However, the effect of skewness is less pronounced compared with the influence of variance.

The majority of the abovementioned papers analyze the influence of learning on risk-taking decisions by comparing decisions between one risky option and a riskless option. However, individuals and companies usually face a broader range of risk management options, and a significant portion of these options does not totally remove the decision-maker's risk exposure. Thus, our paper amends the literature by examining how experiential learning affects the propensity to invest in risk management activities, such as self-insurance (reduction of loss size), self-protection (reduction of the probability of loss), and combined prevention (simultaneous investments in self-insurance and self-protection).

The vast body of literature on prevention decisions and risk aversion based on an expected-utility framework provides different predictions for self-insurance investments compared with self-protection

[1] Burgos (2002) finds for two further learning models that risk-averse behavior is more pronounced in the gain domain. In addition, adaptive learners seem to be prone to the certainty effect and thus choose the less risky of two options more often when it provides a certain payoff.

activities. Ehrlich and Becker (1972) show in their seminal paper that (self-)insurance and self-protection can be substitutes or complements. In addition, the influence of an increase in risk aversion on both prevention technologies has been debated in the literature for decades. Contrary to common intuition, an increase in risk aversion does not necessarily lead to an increase in self-protection, whereas more risk-averse individuals always invest more in self-insurance (Dionne & Eeckhoudt, 1985). Viscusi (1979) examines the influence of learning on risk management decisions in an expected utility setting and argues that self-protection impedes decision-makers' learning about the true probability of loss for a certain risk. In general, a greater precision of the decision-makers' initial assessment of the probability of loss increases the attractiveness of insurance but reduces investments in self-protection.

This research project aims to extend the prevention literature by analyzing the effect of probability of loss and experience on individuals' demand for risk management in an experiential learning model. Additionally, we examine whether experiential learners regard self-insurance and self-protection as complements or substitutes. In contrast to the extensive literature that uses expected utility models to explain prevention decisions, to our knowledge, only very few papers relate experiential learning models and risk management decisions. Experiential learning models require no assumption about individuals' beliefs about loss distributions and the effectiveness of risk management options. These models simply assume that individuals respond to outcomes in prior periods in a predefined way. The lack of descriptive information when making risk management decisions seems to be more in line with real-life decision-making.

One paper that links experiential learning and risk management decisions is Shafran (2011) who conducts several multi-period experiments to analyze how subjects adjust their decisions to invest in self-protection depending on the type of risk and feedback. Subjects are more likely to purchase protection against high-probability, low-consequence losses than against low-probability, high-consequence losses. Although the author provides full information about loss distributions at the beginning of the experiment, subjects respond to feedback.

In addition to Shafran's work, Meyer (2012) examines investments in self-insurance measures over time, depending on loss experience. He uses a simple trial-and-error learning model to predict individuals' investments. In his experiments, subjects underinvest in self-insurance despite frequent feedback and only increase their investments immediately after experiencing losses. However, in line with empirical evidence on the demand for catastrophe insurance, his subjects reduce self-insurance investments again in subsequent periods.

Our paper amends the literature by analyzing both self-protection and self-insurance investments in one simulative study. In addition, we are able to account for the heavily skewed loss distribution because we control for the size of the final outcome. In contrast, Meyer assumes that self-insurance mitigates a loss completely. Thus, the positive reward for a prevention measure is restricted to the occurrence of a disaster. We are able to introduce incomplete self-insurance, and as a result, self-protection investments are not superfluous, as they are in Meyer's model. Hence, we are able to introduce a more differentiated feedback mechanism for investments in loss mitigation.

Model

We develop a basic model of experiential learning to examine individual risk management decisions[2] over time depending on the probability of loss. The decision-maker initially has no information about the payoff distribution and does not know how his risk exposure is exactly affected by each risk management activity. Importantly, we do not impose initial preferences for any risk management option or for the "no prevention" option. We only assume that the individual evaluates a selected option's outcome relative to a predefined reference outcome (aspiration level).

The individual in our model faces a loss with a probability of loss PL in each period.[3] If a loss event occurs, the decision-maker will suffer from a loss l. In each period, the individual can decide to invest in prevention. Subjects can choose between two risk management options. The first technology (self-insurance) reduces the size of a loss when the disaster occurs by a fixed percentage m (e.g., 80%). The second technology (self-protection) does not influence the loss size but reduces the probability of the loss event by φ (e.g., 50%).[4] We assume that both technologies cause fixed costs c_1 and c_2, which hold the expected value (EV) of the individual's final wealth constant. The payoff characteristics of the three basic strategies in our simulation are exemplarily summarized for a baseline probability of loss PL = 0.1 and loss size of 120. We assume that subjects' initial wealth level is 140[5] in Table 1:

No.	Alternatives:	Loss prob.	Loss size	Cost	Expected Value
(1)	No prevention	0.1	120	0	128
(2)	Self-insurance	0.1	**20**	10	128
(3)	Self-protection	**0.05**	120	6	128

Table 1: Decision-maker's Strategy Set

Initially, the probability of choosing one option $\pi_0(a_i)$ is the same for all available options because individuals have no information about the effectiveness of each prevention technology. Decision-makers adjust their initial preferences over time after observing the outcomes of their risk management decisions.

[2] We do not explicitly analyze market insurance because the payoff structure in our model does not differ from the self-insurance case. Hence, our results on self-insurance can be applied to market insurance as well.
[3] We rule out multiple loss events for one individual within one period.
[4] In order to examine the substitutability of both options in the second part of the paper, we assume that both technologies provide only imperfect protection.
[5] We assume 140 as the initial wealth level to ensure positive outcomes and to avoid changes in the payoff domain after deducting risk management costs.

In line with prior stochastic learning models (e.g., Bush & Mosteller, 1955; Luce, 1959; March, 1996; Denrell & March, 2001), we assume that outcomes in prior rounds which exceeded the aspiration level increase the probability of choosing one alternative again, and negative experiences decrease the likelihood that this alternative will be selected in the following periods. Our model is a simplified version of the "experience-weighted attraction" learning model developed by Camerer and Ho (1999), and we essentially adopt the reinforcement comparison model suggested by Sutton and Barto (1998).

The probability of choosing one alternative depends on the relative attractiveness of an option $\boldsymbol{P_t}(a_i)$ compared to the other options in this period. Therefore, the probability of selecting alternative a_i in the t^{th} round is

$$\boldsymbol{\pi_t}(\boldsymbol{a_i}) = \frac{e^{P_t(a_i)}}{\sum_{j=1}^{n} e^{P_t(a_j)}} \qquad \textbf{[1]}$$

Because only the attraction of the selected alternative a_i is adjusted in period t, the probability of choosing alternative a_i is updated directly. The probability of choosing the other options is altered as a consequence of a_i's probability adjustment. As Camerer and Ho (1999) outline, the exponential form is invariant to adding a constant parameter to the attraction levels $P_t(.)$. Thus, negative values for $P_t(.)$ do not cause any problems, and we can use this updating mechanism to analyze our research question, which is mainly concerned with losses.

In contrast with other models of reinforcement learning, we account for the size of the outcome and do not restrict our examination by classifying outcomes as positive or negative. Distinguishing the size of the outcome r_t is important for analyzing risk management decisions because individuals often face loss distributions that are negatively skewed, with extreme losses occurring with a very low probability of loss. If we did not control for the size of the outcome, incomplete self-insurance investments would be evaluated as poorly as no prevention in the case of a loss. Because risk management measures only pay out in the rare case of a loss event, it is important to capture the magnitude of a loss when prevention has not been undertaken.

Both individuals and organizations evaluate the outcomes of their actions in comparison to reference points. Decision-makers' willingness to take risks depends on whether their payoff is above or below their aspiration level (Kahneman & Tversky, 1979; March, 1988; March & Shapira, 1992). In order to evaluate the outcome of each risk management option, we define a reference outcome ($\bar{r}_t$). The outcome of the chosen alternative r_t in one period relative to the reference outcome ($\bar{r}_t$) determines the attractiveness of this alternative in the next round:

$$\boldsymbol{P_{t+1}}(\boldsymbol{\alpha_i}) = \boldsymbol{P_t} + \boldsymbol{\beta}[\boldsymbol{r_t} - \bar{\boldsymbol{r}}_{\boldsymbol{t}}] \qquad \textbf{[2]}$$

$$\boldsymbol{\beta} > \boldsymbol{0}$$

The current attractiveness level P_t is updated by the outcome difference, which is weighted by a positive parameter β. β determines the speed of learning and thus influences how strongly the attractiveness of a chosen alternative is altered by current outcomes. We examine how the size of β alters decision-makers' willingness to undertake risk management in the section "Sensitivity Analysis".

We assume that the decision-maker can only learn about the quality of one option if the respective alternative is actually chosen. Therefore, the attraction of non-selected alternatives is not affected by the updating process. In contrast, Camerer and Ho (1999) and Meyer (2012) assume that subjects can learn about the quality of an option by directly observing realized payoffs and by envisioning foregone payoffs or losses that could have occurred. For the sake of simplicity, we restrict our analysis to the direct learning channel because it is unclear what determines the strength and direction of indirect learning.

The reference outcome $\bar{r}_t$, which functions as a benchmark for all options, is also updated in each round. In general, the reference outcome is an incremental average of all selected options. However, the updates of the reference outcome take place after the attraction level update in (2).

$$\bar{r}_{t+1} = \bar{r}_t + \alpha[r_t - \bar{r}_t] \quad \textbf{[3]}$$

$$0 \leq \alpha \leq 1$$

The sensitivity of the updating process is determined by α. α is bounded between 0 and 1 and is usually assumed to be smaller than the learning parameter β. In addition, the starting value of the reference outcome $\bar{r}_0$ influences the decision-maker's tendency to explore other options. For extreme starting values, the learning algorithm may break down. In general, the reference outcome should converge to the expected value rather quickly, according to prior literature (Jaspersen & Peter, 2014). To examine the impact of the initial value of the reference outcome on our results, we also conduct a sensitivity analysis. We use the expected value of our options as the starting value of the reference outcome in our baseline setting.

Distributional Properties of Risk Management Options

Denrell (2007) shows that learning leads to risk-averse behavior when distributions are symmetric but may also result in risk-seeking behavior if distributions are negatively skewed. Because we are interested in how learning influences the take-up of different risk management options, we have to examine the first, second and third moments of their payoff distributions. Table 2 shows the basic distributional properties of our risk management options at a probability of loss of 0.1:

	No prevention	Self-insurance	Self-protection
Payoff No Loss	140	130	134
Payoff Loss	20	110	14
Probability of loss	0.1	0.1	0.05
Mean	128	128	128
Variance	1296	36	684
Standard Deviation	36	6	26.15
Skewness	-2.66	-2.66	-4.13
Mode	140	130	134

Table 2: Distributional Properties of Risk Management Options

In our baseline case, all options provide the same mean, as we assume actuarially fair costs. Investments in self-insurance reduce the payoff distribution's variance and transfer risk from the tails to the center of the distribution. Thus, the self-insurance option in our model leads to a mean-preserving contraction of the distribution. As a consequence, Briys and Schlesinger (1990) show that fairly priced self-insurance activities second-order stochastically dominate the "no prevention"-option, and risk-averse decision-makers should strictly prefer the investment. Additionally, the distribution's skewness is not affected by self-insurance. According to the literature, learning results in risk-averse behavior for symmetric distributions. Thus, we assume that learners focus on the variance of both distributions if two investments do not differ with regard to their degrees of skewness.

In contrast to self-insurance, investments in self-protection usually alter both variance and skewness of the distribution. If we assume that the probability of loss is lower than 0.5, self-protection lowers the variance, but the distribution is more skewed. Briys and Schlesinger (1990) describe the impact of self-protection as a mean-preserving contraction combined with a mean-preserving spread when self-protection costs are actuarially fair. Thus, a clear prediction about the preferences of risk-averse decision-makers is not possible because subjects have to trade-off the change in variance with the change in skewness (see, e.g., Chiu, 2005). In general, learning should result in risk-averse and skewness-averse behavior (Jaspersen & Peter, 2014). Hence, the influence of learning on self-protection take-up rates should be less pronounced than the effect of learning on self-insurance investments.

In the majority of cases, adaptive learning identifies the payoff-maximizing strategy. However, Erev and Barron (2005) describe important paradigms when adaptation deviates from payoff maximization. For instance, if one alternative provides the best outcome most of the time but has a lower expected value compared with another option, decision-makers may still prefer the first alternative because they behave as if they underweight rare outcomes. Our alternatives provide the same expected value, but the "no prevention"-option still provides the highest mode of all available options. When the probability of loss decreases, positive feedback for investments in risk management is rare, and the importance of the "no-loss" outcome increases.

In addition, a change in the probability of loss also influences the differences of the distributional properties between the available options. Figure 1 and Figure 2 show how the difference of the distributional properties between the "no prevention"-option and the "self-insurance"-option and the "self-protection"-option, respectively, changes depending on the probability of loss.

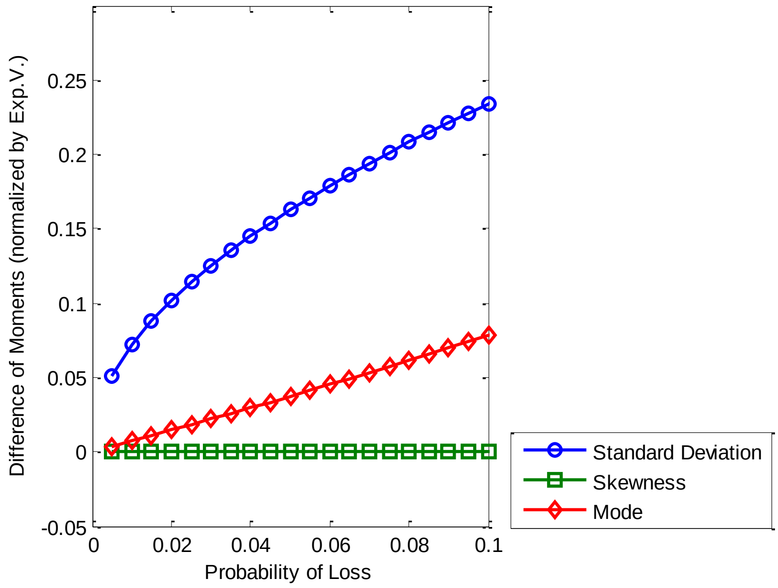

Figure 1: No Prevention and Self-insurance: Difference of Distributional Properties[6]

Because there is no difference in skewness between the "no prevention" and the "self-insurance" options, a change in the probability of loss does not affect the preference ranking of either option with respect to skewness. However, both the difference between the standard deviation and the difference between the modes decrease when we lower the probability of loss. Thus, we observe two counteracting effects of the probability change on the relative attractiveness of self-insurance.[7]

Figure 2 shows the impact of lowering the probability of loss on the difference in standard deviation, skewness and mode between the "no prevention" and "self-protection"-options. We receive the same picture for the standard deviation and mode, but additionally, we observe an increase in the difference of the skewness coefficients when the probability of loss decreases. Hence, changing the probability of loss implies another trade-off in the case of no prevention vs. self-protection because decision-makers in experiential learning models are both variance and skewness averse.

[6] We subtract self-insurance standard deviation, skewness and mode from the "no prevention" option's respective distributional properties.
[7] We divide standard deviation, skewness and mode by the expected value of all options at the respective probability of loss to ensure the comparability of the differences between both options at different probability levels.

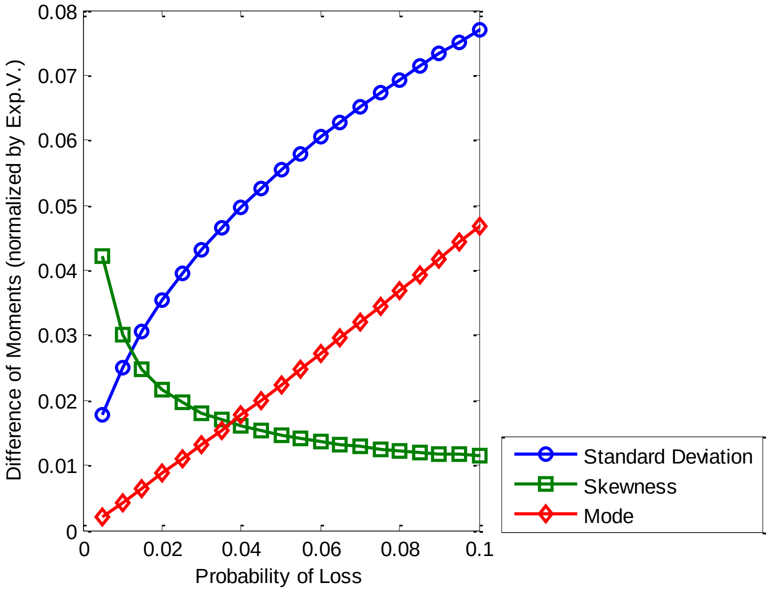

Figure 2: No Prevention and Self-protection: Difference of Distributional Properties[8]

Simulation

This section provides an overview of the different simulation runs that we conduct to answer our research questions. As the first step, we examine the propensity to invest in self-insurance or self-protection in two separate simulation runs in which the decision-maker can only choose between no prevention and one prevention technology, respectively. We vary the probability of loss between 0.005 and 0.2 in incremental steps of 0.005. Table 3 summarizes all parameters we intend to use in our simulation models and their values in the baseline case.

Parameter	Description	Default Value
t	number of rounds	50
i	number of individuals	10,000
$\boldsymbol{\alpha}$	adjustment of reference outcome	0.1
$\boldsymbol{\beta}$	adjustment of option's attractiveness	0.6
$\boldsymbol{\bar{r}_0}$	reference outcome in period 0	Expected Value
$\boldsymbol{\pi_0(a_i)}$	probability of choosing option i in period 0	$\frac{1}{\text{Number of Options}}$
$\boldsymbol{P_0(a_i)}$	attractiveness of option i in period 0	Expected Value

Table 3: Summary of Simulation Parameters

We choose a fairly low adjustment parameter α for the reference outcome because we assume that decision-makers update their beliefs about the reference value only very slowly. In contrast, we set a

[8] We subtract self-protection standard deviation, skewness and mode from the "no prevention" option's respective distributional properties.

higher value for β, the adjustment parameter for the option's attractiveness, as the occurrence of a loss event will change the evaluation of the selected option more quickly. To analyze risk management decisions over time depending on loss experience, we simulate decisions of 10,000 individuals over 50 time periods.

In addition to developing a sound descriptive model for analyzing risk management decisions over time, our second goal is to contribute to the discussion of whether self-protection and self-insurance are substitutes or complements. Our model aims to shed light on how the investment in one prevention technology influences the likelihood or extent of learning opportunities about the other prevention measure. Thus, in a third simulation run subjects make both risk management decisions simultaneously. Four different strategies are available: no prevention, self-protection or self-insurance or both prevention options simultaneously. However, the probability of selecting each option is updated separately. Hence, the subject makes two distinct decisions whether to invest in self-insurance or not and whether to invest in self-protection or not. The decision-maker receives feedback about the quality of an option only if he invests in the respective option. However, if both options are selected, feedback is affected by the joint outcome.

After examining experiential learners combined risk management decisions at fair preventions costs, we introduce a loading factor and show how the probability of investing in both risk management options is affected by varying the costs of one option. The loading factor is altered in incremental steps of 0.1 between 0.1 and 5. As a result, the attractiveness of the respective risk management option in terms of expected value is modified as well. Because reinforcement learning models respond heavily to differences in the expected value between options, we expect a decline in the probability of selecting a risk management option with positive loading (= loading factor >1). However, the effect on the alternative risk management option is less clear.

To answer this question, we conduct two additional simulative studies in which the decision-maker can choose between no prevention and both risk management options. We introduce the loading factor for only one option per simulation run and analyze how both take-up rates vary depending on the size of the loading. If we increase the cost of option 1 and if take-up rates of option 2 increase, we regard this as evidence of a substitutive relationship. In contrast, if the take-up rates of option 2 decrease when increasing cost of option 1, we consider both options as complements.

Results

Experience and Probability of Loss

In order to examine the impact of learning on risk management decisions, we conduct two separate simulation runs. Firstly, we test how the propensity to invest in self-insurance changes over 50 rounds. Figure 3 shows our results for different levels of probability of loss:

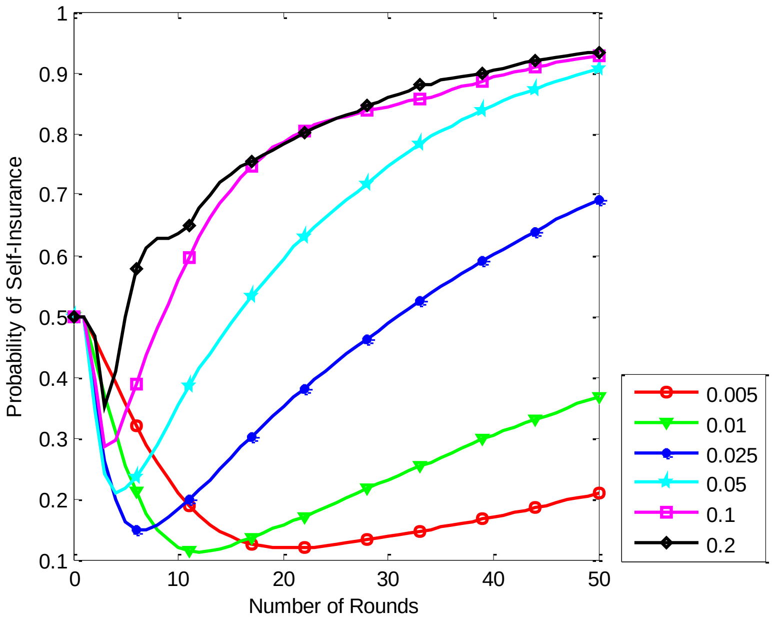

Figure 3: Probability to invest in self-insurance depending on probability of loss

We compute the average probability of investing in self-insurance for 10,000 individuals over time. The mean of the probability of investing in self-insurance decreases in the first periods but rises steadily after the decision-makers gain more experience. The decline at the beginning may be driven by the higher pay-off of the “no-prevention”-option in the “no loss”-state (equivalent to the mode) relative to the aspiration level compared to the “self-insurance”-option. The positive reinforcement for investments in the “no-prevention”-option is stronger and only after the occurrence of loss events the likelihood of investing in self-insurance picks up. Thus, we observe the minimum for lower probabilities of loss in later rounds. For instance, the probability of investing in self-insurance at PL = 0.005 decreases approximately until t=20. However, all take-up rates of self-insurance increase in the long run. In line with prior literature on learning and risk taking, we find that more learners act risk averse over time (March, 1996; Denrell & March, 2001; Denrell, 2007).

To show the effect of the probability of loss on risk taking behavior in a model of reinforcement learning, we depict the average probability of investing in self-insurance after 50 periods for different levels of probability of loss in Figure 4. We vary the probability of loss incrementally from PL = 0.005 to PL = 0.2 in steps of 0.005:

In general, we receive a fairly uniform relation between self-insurance take-up rates and probability of loss. The probability of investing in self-insurance rises heavily with increasing probabilities of loss and stabilizes at take-up rates of about 0.9.

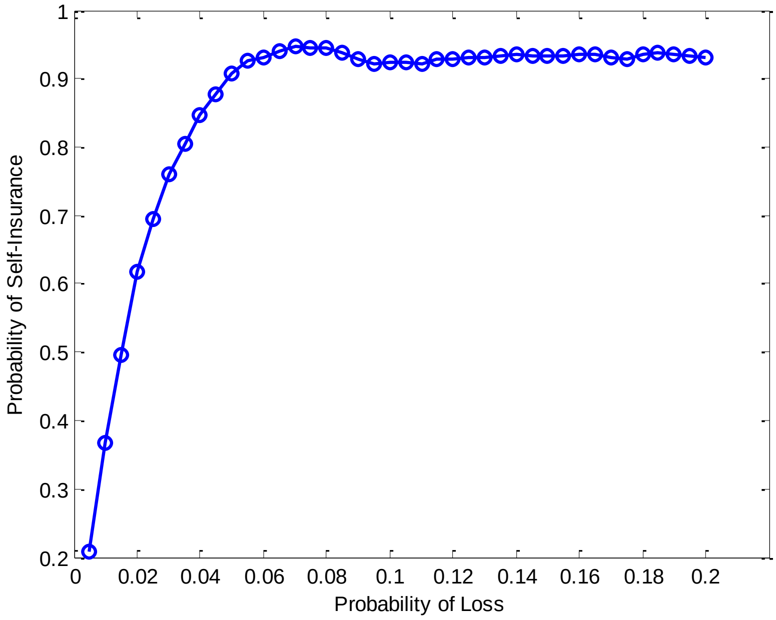

Figure 4: Probability of Investing in Self-Insurance at t=50

Our second simulation run examines the take-up rates of the "self-protection"-option among 10,000 individuals over time. Figure 5 summarizes our findings for different levels of probability of loss. Interestingly, we receive a fairly different picture. First of all, the shape of the "learning curve" is different for large to medium initial probabilities[9] (0.05-0.2). We observe only a short increase in take-up rates after the initial decline in the first rounds. Take-up rates are fairly stable in the long run for these probability levels. In contrast, the development of take-up rates is similar to the self-insurance case for small to medium probabilities of loss (0.005-0.04). In general, the influence of experiential learning seems to be weaker as the highest take-up rates are about 0.7 in comparison to self-insurance take-up rates of 0.9.

The assumption of a non-linear effect of the probability of loss on self-protection take-up rates is supported after depicting investment probabilities after 50 time periods depending on the probability of loss. Figure 6 shows that the probability of investing in self-protection increases with rising probability of loss up to a maximum of 0.68 at about PL = 0.04. Afterwards the take-up rate drops quickly to rates of about 0.6. This observation is in strong contrast to the predominantly positive relationship between self-insurance take-up rates and probability of loss. We conclude that the optimal trade-off between variance reduction and skewness increase when investing in self-protection seems to be at medium size probabilities of loss.

To summarize, take-up rates of both risk management options are low for small probabilities of loss which is in line with empirical observations of underinvestment / underinsurance against low-probability risks. However, take-up rates differ for higher loss probabilities. Self-insurance is adopted for high probabilities of loss but self-protection seems to be most popular at medium levels of probability of loss.

[9] We refer to probabilities of loss before the implementation of self-protection if not stated otherwise.

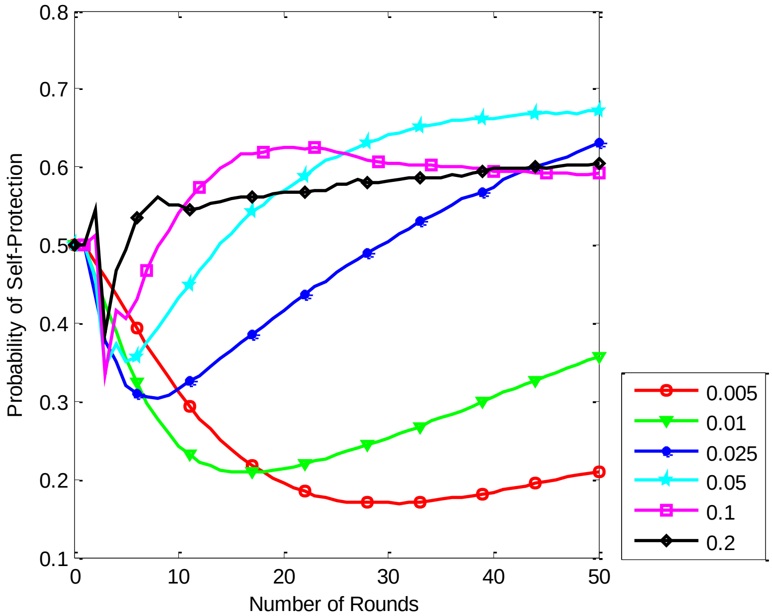

Figure 5: Probability to invest in self-protection depending on probability of loss

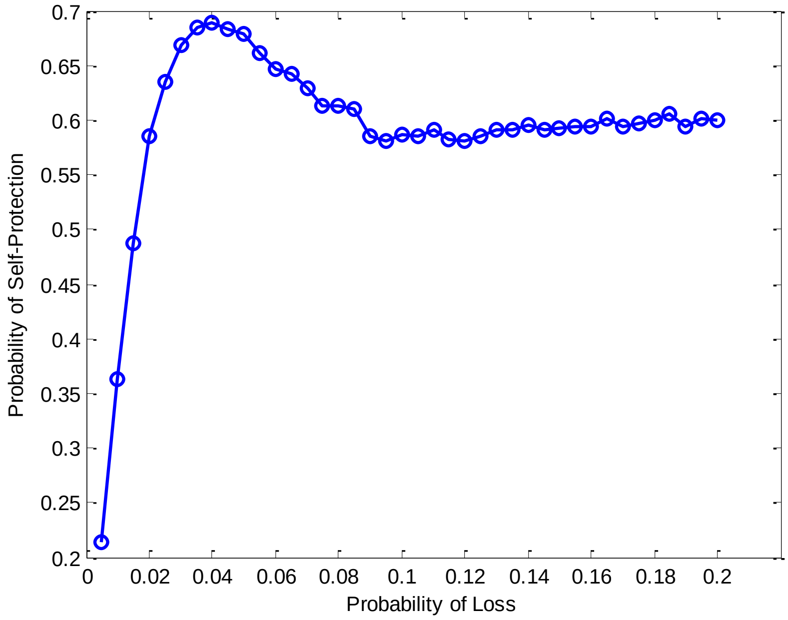

Figure 6: Probability of Investing in self-protection at t=50

Combined Prevention

After examining the effect of probability of loss and experience on experiential learners' single risk management decisions, we extend the choice set in one simulation run to analyze interdependencies between both options. In this more realistic setting, decision-makers can choose between no prevention, self-insurance, self-protection and both options at the same time. If we compare take-up rates of both risk management options over time at different levels of probability of loss in the combined prevention setting, we receive similar results as in the separate simulation runs. The development of the take-up rates of both options over time is shown for selected levels of probability of loss in Figure 12 in the appendix. At low probabilities of loss (0.005 to 0.025) take-up rates are pretty much alike. However, for higher probabilities of loss we observe, as described before, a diverging development. As in the separate simulation runs, self-protection take-up rates level-off whereas self-insurance take-up rates increase during the entire learning period. The outcome of the learning process is summarized in Figure 7. Take-up rates of both options after 50 periods rise simultaneously when we increase the probabilities of loss up to about 0.05. Afterwards self-insurance take-up rates remain at this high level, whereas self-protection take-up rates are lower at high probabilities of loss.

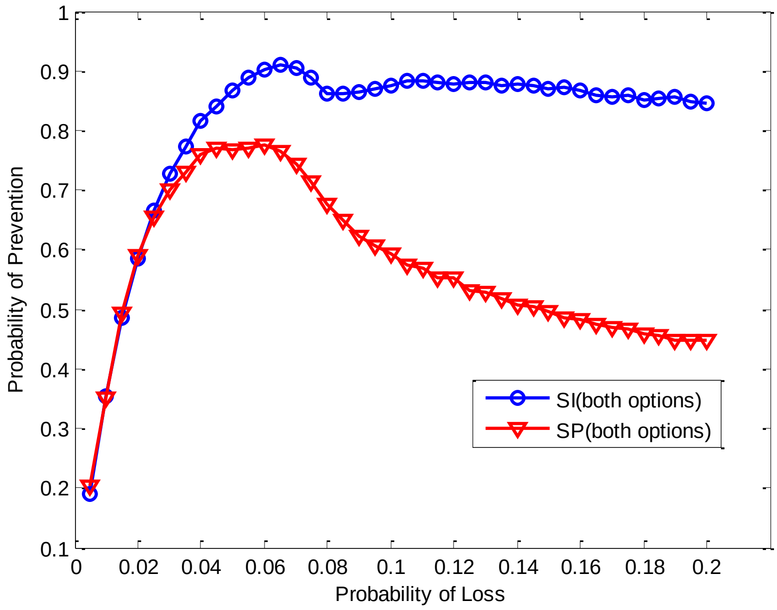

Figure 7: Combined Prevention: Probability of Investing in self-insurance and self-protection at t=50

To highlight differences between the results of the simulation runs with one risk management option and the simulation offering both options, we compare in Figure 8 and Figure 9 take-up rates in the “separate” prevention simulation with the “combined” prevention simulation.

We receive overall lower self-insurance take-up rates in the “combined” simulation run which could indicate a substitutive relationship between both options. However, the relation between the size of the probability of loss and the self-insurance take-up rates does not change. Higher probabilities of loss still result in higher self-insurance take-up rates.

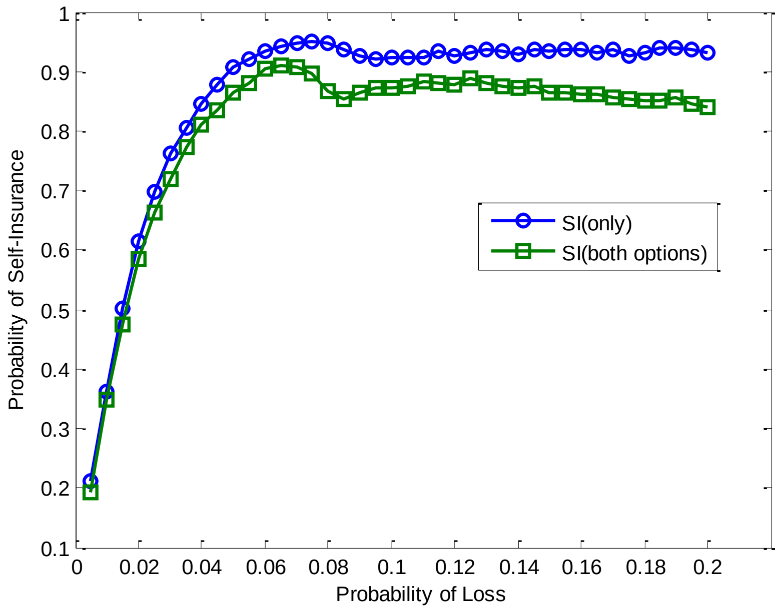

Figure 8: Probability of Investing in Self-Insurance (Single vs. Combined Prevention) at t=50

Figure 9 emphasizes differences between both simulation runs for self-protection take-up rates. At very low levels of probability of loss, the propensity to invest in self-protection is almost the same in both settings. In contrast, we observe higher self-protection take-up rates at medium size probabilities of loss in the "combined" prevention case. This may be the result of positive spill-over effects of self-insurance investments because the drawback of a self-protection investment of lower pay-offs in the loss state compared with the "no prevention"-option is mitigated. However, at probabilities of loss above 0.1, self-protection take-up rates are lower in the "combined" prevention case compared with the separate self-protection simulation.

Lower take-up rates in the "combined" prevention setting may be driven by the fact that investing in both self-insurance and self-protection at the same time results in a lower expected value compared with the "no prevention"-option and the separate investment in one risk management option. The cost of prevention in our simulation models is based on the reduction of expected loss by investing in the respective risk management option. Cost of one risk management option does not decrease if the decision-maker invests in the second option because we assume two distinct investment decisions.[10] Investing in both risk management options at the same time also results in an increase in skewness compared with the "no prevention"-option and compared with the stand-alone "self-insurance"-option. However, payoffs of the combined prevention option exhibit the lowest variation of all available options. A comparison of the distributional properties of combined prevention with the other three options is shown in Figure 13-15 in the appendix.

[10] If we assume that the cost of one option also depends on investments in the other option, we would introduce a complementary relationship between both risk management options directly.

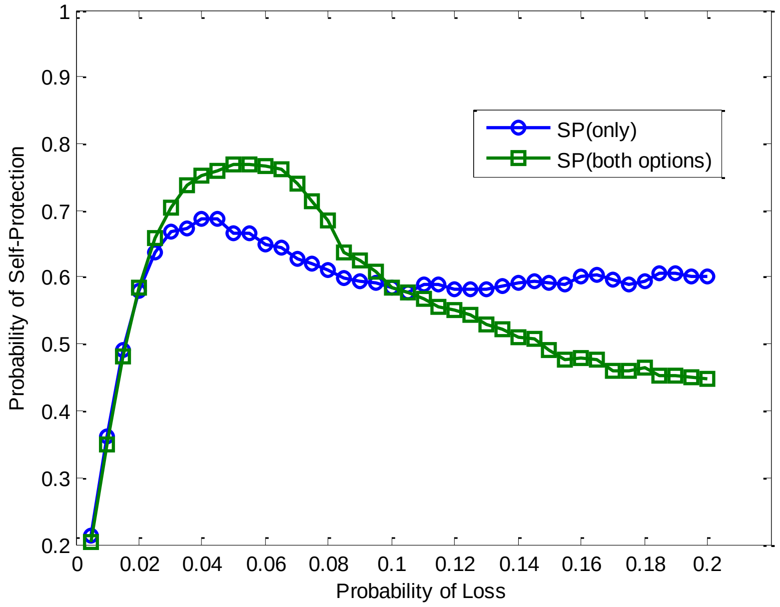

Figure 9: Probability of Investing in Self-Protection (Single vs. Combined Prevention) at t=50

Substitutes or Complements

To examine the interdependencies between self-insurance and self-protection decisions comprehensively, we vary the costs of one risk management option and show how the take-up rates for both options are affected. The loading factor in our simulative study is adjusted incrementally between 0.1 and 5 with steps of 0.1. Figure 10 contains the results for both options when the loading factor of the self-insurance costs is altered. The effect of the loading depends on the probability of loss. Learners' self-insurance take-up rates are price sensitive for probabilities of loss between 0.05 and 0.2. Interestingly, there is no noteworthy effect of price changes on self-insurance take-up rates for lower probabilities of loss. Since reinforcement learning models typically respond to differences in expected value, the effect of the loading factor on the expected value might be too small when the probabilities of loss are very low. To determine the substitutability between both risk management options, we examine the effect of changing self-insurance costs on the probability of investing in self-protection. For most probability levels, we do not observe a strong effect on self-protection take-up rates when we vary the self-insurance loading. Medium probability levels of 0.05 might be an exception because the probability of investing in self-protection varies in a range of 20 percent. However, the effect is inconclusive because we observe an increase in the take-up rate of self-protection when the loading factor increases from 0.1 to 1. This would indicate that both options were substitutes. Interestingly, the probability of investing in self-protection decreases again for loading factors in the range of 1 to 5. At probabilities of loss below 0.05 we observe some cross-price effects which indicate a weak substitutive relationship. The results for both risk management options for different self-protection loading factors at various probabilities of loss are summarized in Figure 11..

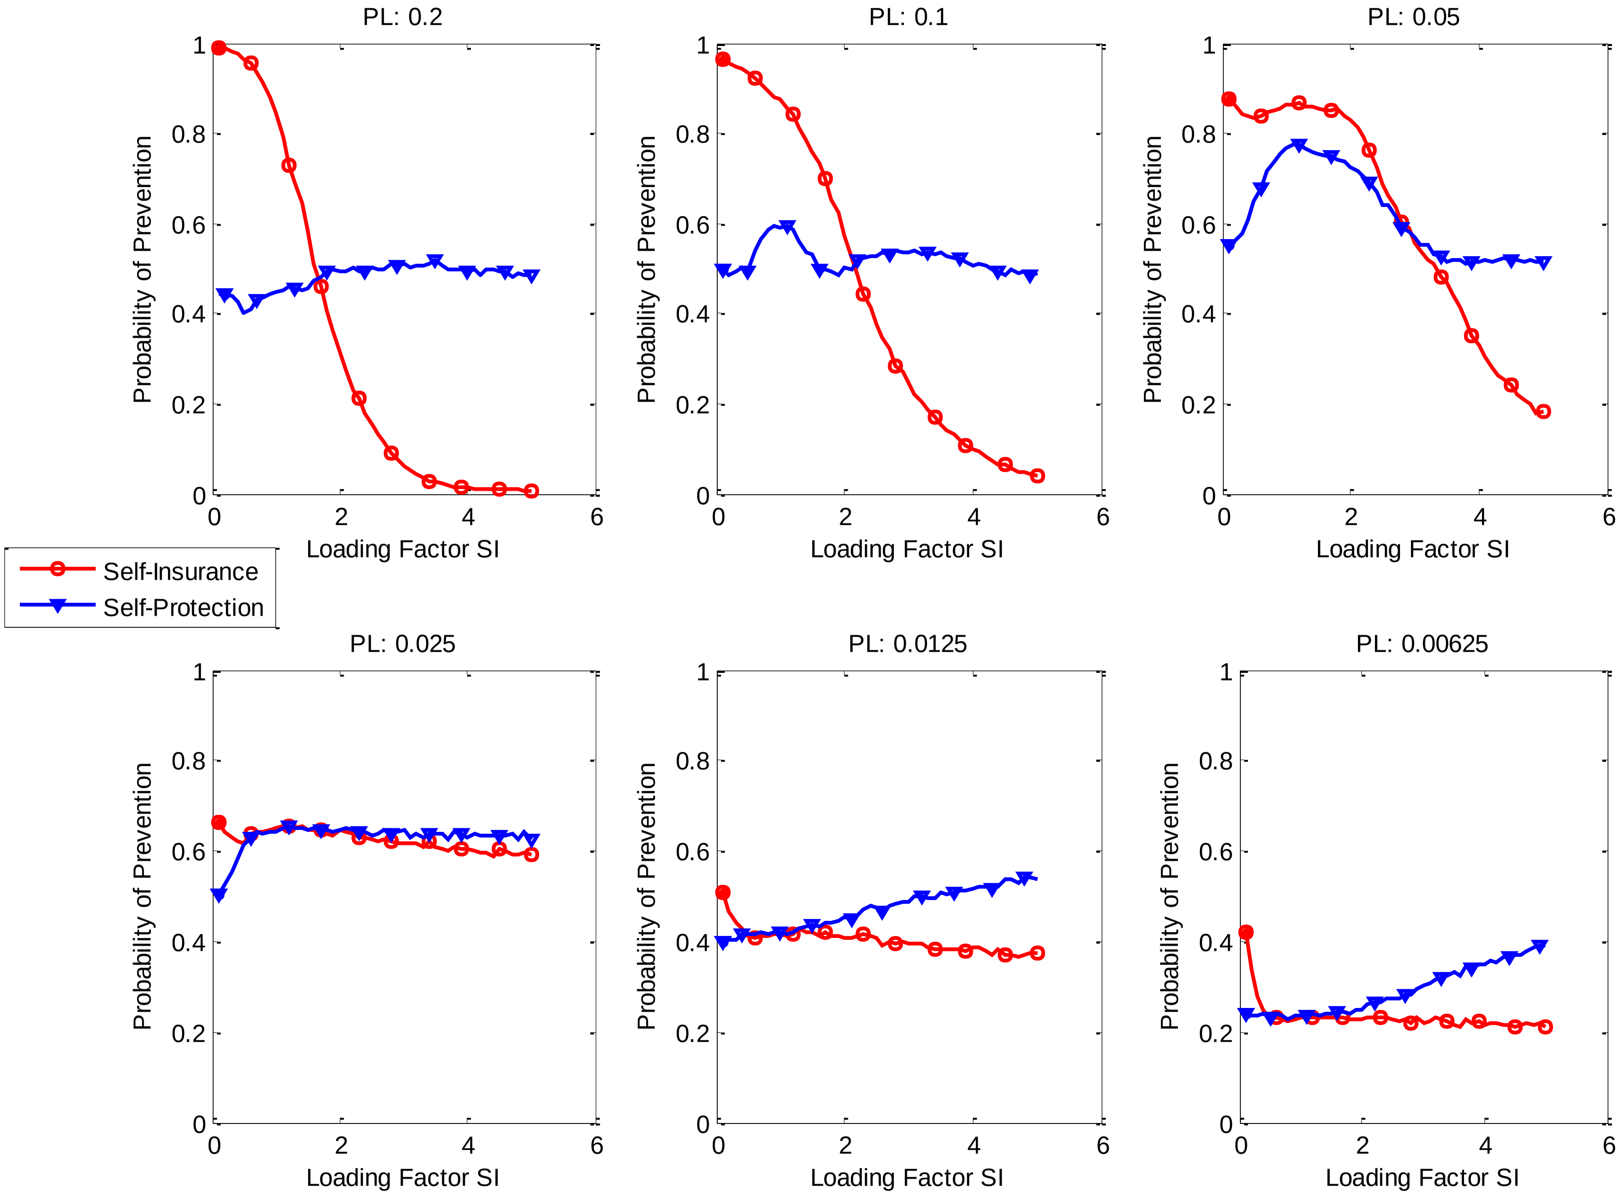

Figure 10: Prevention Investments at t= 50 depending on PL and SP-Loading Factor

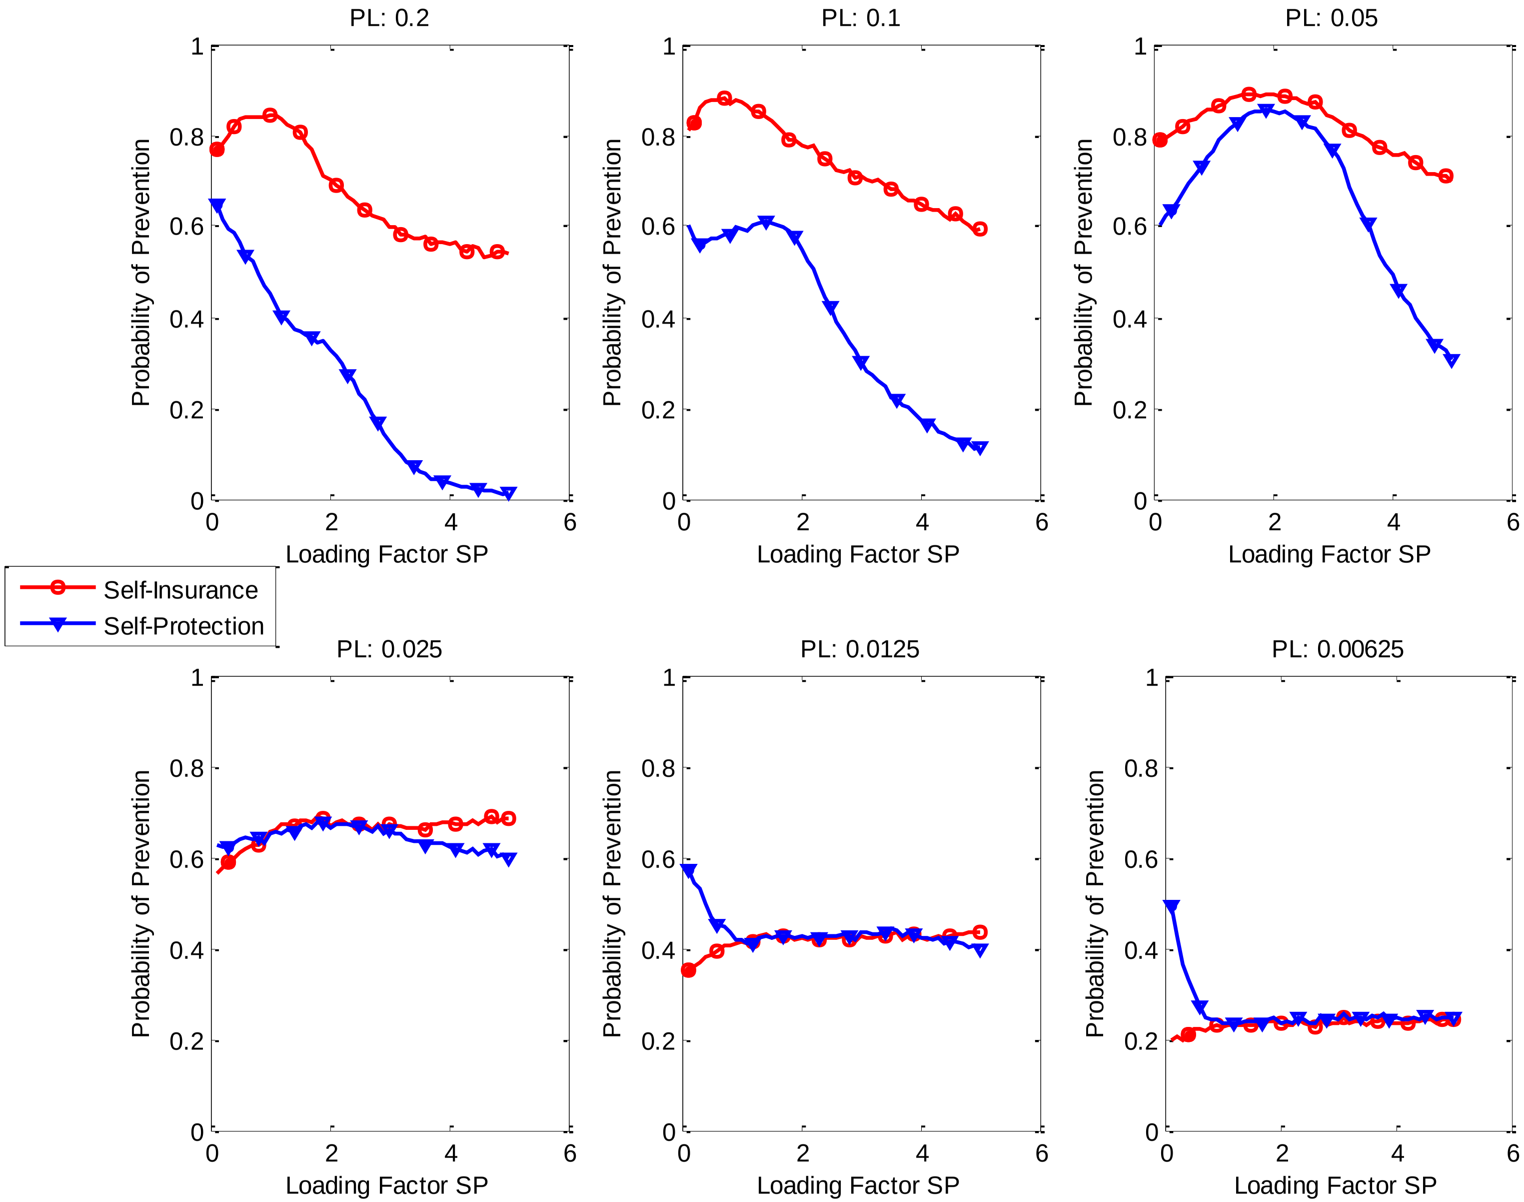

Figure 11: Prevention Investments at t= 50 depending on PL and SP-Loading Factor

Learners' take-up rates for self-protection seem to be price-sensitive only for high probabilities of loss (0.1 – 0.2). While we do not observe a price effect for very low probabilities of loss, the effect of price changes at medium probabilities of loss is counterintuitive. For PL = 0.05 the probability of investing in self-protection reaches its maximum at loading factors between 1 and 2 and drops at higher and lower loading factors. We do not observe a remarkable cross price effect on self-insurance investments for small to medium probabilities of loss. However, self-insurance take-up rates are affected by self-protection price changes for probabilities of loss between 0.05 and 0.2 indicating a complementary relationship. In general, take-up rates of both risk management options seem to be price-sensitive for most probabilities of loss in our learning model. At some probabilities of loss, the two options behave in a complementary way when varying the self-protection loading factor. In contrast, we receive a weak substitutive relationship at low probabilities of loss when varying the self-insurance loading factor. Because self-insurance investments do not affect the skewness of the distribution but reduce the variance, the learning algorithm seems to favor this risk management option. In contrast, investments in self-protection imply a trade-off between an increase in skewness and a reduction of variance which results in lower take-up rates overall.

Sensitivity Analysis

We conduct an extensive sensitivity analysis to examine the effect of the size of the updating parameter α and β and of the initial aspiration level (reference outcome). In addition, we analyze long-term learning effects by extending the time period to 100, 500, and 1000 rounds, respectively. Finally, we compare experiential learners risk management decisions against low probability, high consequence risks with high probability, low consequence risks. Results of our sensitivity analysis are shown in the appendix.

Aspiration Level Updating Parameter α

We set α = 0.1 in our baseline setting because we assume that individuals' aspirations are fairly stable and are updated slowly. We vary α incrementally from 0.1 to 1.0 while holding β, the updating parameter for the attractiveness of the available options, constant at 0.6. As long as individuals update their aspiration level slower than their preferences for the various options, we do not observe a strong effect when varying α as shown in Figure 16**Fehler! Verweisquelle konnte nicht gefunden werden.** and Figure 17**Fehler! Verweisquelle konnte nicht gefunden werden.**. However, take-up rates of both risk management options decrease for large values of α. If individuals' aspiration levels are solely determined by the most recent outcome (i.d., α = 1.0), we observe a strong positive relationship between take-up rates of both risk management options and probabilities of loss.

Preference Updating Parameter β

Our results seem to be less sensitive to variations of the preference updating parameter. We vary β between 0.1 and 1.5 and detect almost no effect on both risk-management options (see Figure 18**Fehler! Verweisquelle konnte nicht gefunden werden.** and Figure 19).

Initial Aspiration Level

The initial aspiration level influences the evaluation of the individuals' first draws and thus, could determine future sampling. Figure 20 and Figure 21 in the appendix depict the average probability of choosing the "self-insurance"-option and "self-protection"-option after 50 time periods for several probabilities of loss levels depending on the initial aspiration level. We vary the starting value of the aspiration level in percent of the expected value of the available options. If the initial value deviates heavily from the expected value (10 % to 75 % of the expected value) both risk management options, and the "no prevention"-option receive positive feedback because their outcome exceed the aspirations significantly. Due to the small aspiration updating parameter the aspiration level adjusts slowly and the learning algorithm breaks down. As a result take-up rates are about 0.5 which is equivalent to random guessing. We receive the highest take-up rates for initial aspiration levels close to the expected value which facilitates the most efficient learning process. At very low probability of loss levels (PL = 0.005), we detect a different pattern. In general, probabilities of investing in prevention are lower and the lowest take-up rates are observed at initial aspirations levels around the expected value.

Learning Periods

For our main analysis, we use a learning period of 50 rounds because we assume annual probabilities of loss and thus, longer time periods seem to have very limited implications for real-life risk management decisions. In our sensitivity analysis, we extend the learning period to 100, 500 and 1,000 rounds. The effect of learning on self-insurance investment is more pronounced for longer time-periods (see Figure 22-24). After 500 or 1,000 periods take-up rates for self-insurance investments are above 0.9 for all probabilities of loss. The effect of learning is weaker for self-protection investments. Take-up rates do not increase further for most probabilities of loss after about 100 periods of learning and remain stable over time (see Figure 25-27). Interestingly, after 1,000 rounds we detect the highest take-up rates when the probability of loss is very low. Due to very low risk management costs at low probabilities of loss in our model both options, self-protection and no prevention, provide almost equivalent payoffs. However, the reduction of the probability of loss pays out in the long run.

Low Probability, High Consequence Risks (LPHC) versus High Probability, Low Consequence Risk (HPLC)

In the result section, we compared take-up rates of self-insurance and self-protection depending on the probability of loss and amount of experience. When varying the size of the probability of loss, we also alter the expected value of the available options. As the expected value of the "no prevention"-option and the risk management option in one simulation run was always the same, a clear comparison of the attractiveness of a risk management option relative to the “no prevention”-option and the other risk management option at different levels of probability of loss was possible. However, in the literature risk management decisions against low probability, high consequence risks are often compared with risk management decisions against high probability, low consequence risks (e.g., McClelland et al., 1993; Laury et al., 2009; Shafran, 2011). Thus, we examine experiential learners’ take-up rates of self-insurance and self-protection, respectively, when varying both the probability of loss and the loss size but holding the expected value constant. The findings are summarized in the appendix in Figure 28 and Figure 29. Our results are predominantly in line with prior findings. Both risk management options exhibit the lowest take-up rates for small probabilities of loss combined with large losses and take-up rates are higher at higher probabilities of loss. However, both risk management options have the highest take-up rates at medium size probabilities of loss and consequently at medium loss size. The difference in take-up rates between medium and high probabilities of loss is considerably smaller for self-insurance investments compared with self-protection. As the expected value is the same for both LPHC and HPLC risks, we conclude that the loss frequency seems to be a more crucial determinant for risk management take-up rates than loss size. Interestingly, Ganderton et al. (2000) find in an experimental study in contrast to Expected Utility Theory and in line with our model that subjects’ decisions exhibit a greater sensitivity to the probability of loss than to the loss size.

Discussion and Conclusion

In this paper, we develop a novel approach to examine risk management decisions when individuals cannot base their decisions on descriptive information but have to rely on prior experience. Although our model relies only on a few basic assumptions, it is able to account for a wide range of human behavior. Besides decision-makers general tendency to act risk averse, our results also capture individuals’ preference for insuring high probability, low consequence risks instead of low probability, high consequence risks.

The learning process and the outcome of our simulative studies provide valuable insights into risk management decisions in a dynamic setting. The initial decrease in take-up rates in both risk management options can be attributed to weaker positive reinforcement in comparison to the “no-prevention”-option in the first rounds because losses are rare events in our model ($PL <= 0.2$) and the starting value of subjects’ aspiration level equals the expected value of all options. Learning or experiencing losses have very clear implications for the propensity to invest in self-insurance. Independent of the probability of loss take-up rates for self-insurance increase over time. Facing a loss event without insurance results in strong deviations from the aspiration level and consequently in a strong decline in the probability of choosing the “no-insurance”-option. The continuous increase in self-insurance take-up rates over time indicates that learners dislike rare but large negative deviations from the aspiration level more than they favor outcomes which are above the aspiration level most of the time. Smaller probabilities of loss result in lower probabilities of investing in self-insurance after

50 rounds. Because negative feedback for “no-prevention” decisions is rare when the probability of loss is low, individuals might refrain from investing in costly self-insurance. In addition, the higher mode of the “no prevention”-option seems to be more attractive than the risk reducing effect of the self-insurance investment.

Self-protection lowers the probability of loss but also causes worse outcomes when a loss occurs due to the costs of prevention. Thus, deviations from the aspiration level are even more severe when subjects have invested in self-protection before compared with the “no-prevention”-option. This might explain the lower take-up rates for self-protection after 50 periods for large probabilities of loss. If the occurrence of a loss event is still fairly likely after investing in self-protection, subjects experience these more severe negative deviations from the aspiration level quite often. Similarly, take-up rates are low for small probabilities of loss. If the probability of loss is already very small, the additional risk reducing effect of self-protection investments might be negligible for reasonable time horizons. As we observe the largest take-up rates for self-protection at PL= 0.04, we conclude that the trade-off between reduction of variance and increase in skewness is optimal at these probability levels (c.f. Figure 2).

Our findings on the substitutability or complementarity of both risk management options are less conclusive. If the decision-maker can invest in both risk management options at the same time, we observe lower take-up rates of self-insurance compared to the single choice setting. In contrast self-protection take-up rates are higher at medium probabilities of loss in the combined prevention setting. However, at high probabilities of loss self-protection seems to be less attractive than in the single choice setting. To highlight the price and cross-price effects in experiential learning models, we varied the cost of both risk management options. In general, the effect of price changes on the difference in the expected value seem to be too small for very low probabilities of loss to have impact on the risk management take-up rates in our simulative study. We observe a direct price effect on risk management investments for larger probabilities of loss. However, similar to expected utility models we receive only ambiguous results with regard to the complementary or substitutive relationships of both risk management options. Depending on the level of probability of loss, self-insurance and self-protection behave as if they were complements or substitutes.

As our simulative study is in line with empirical evidence on underinvestment against low probability, high consequence risks (Kunreuther, 1996; Kunreuther et al., 1998; Schade et al., 2012), we derive some recommendations for overcoming the underinsurance problem. Our results show that the price of risk management investments seems to be of minor importance at very low probabilities of loss. Even if the price was heavily subsidized (e.g., 0.1 * expected cost), we observe a fairly weak effect on take-up rates at small probabilities of loss. The National Flood Insurance Program in the U.S. seems to face similar problems. Despite of significant subsidization of insurance premiums in the past, the demand for flood insurance remained low and is in general considered as rather price-inelastic (Browne &

Hoyt, 2000; Kriesel & Landry, 2004). Thus, offering subsidized insurance contracts might be an inappropriate measure to increase insurance take-up for low probability risks. Connor (1996) highlights that individuals might consider insurance as an investment, similar to our model, when they make their risk management decisions. Thus, they expect to receive a monetary return from their insurance policies. To encourage insurance demand for low probability risks, insurers could provide (partial) premium reimbursement when no losses occur. Slovic et al. (1977) show the positive effect of premium refunds in an experimental setting. In our model, reimbursements will increase the payoff of the insurance option in the "no-loss" case and should increase the positive enforcement for the investment.

Subjects in our model start with zero experience about the effectiveness of different risk management investments. Hence, it could be used to examine the appropriateness of competing options for developing risk management strategies in new markets or against new risks. For instance, microinsurance is a fairly new form of insuring low-income people in developing countries. As formal insurance contracts were not accessible in these markets, the majority of people have only limited experience with insurance products (Cohen & Sebstad, 2005; Roth et al., 2007). Thus, our model can help to predict the development of new insurance markets where policyholders have to gain experience about the benefits of formal insurance products over time. Our simulative studies can provide a first indication for which policy type, or risk management option in general, would lead to the highest take-up rates or fastest dissemination in a new market.

References

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16*(3), 215-233.

Briys, E., & Schlesinger, H. (1990). Risk aversion and the propensities for self-insurance and self-protection. *Southern Economic Journal, 57*(2), 458-467.

Browne, M. J., & Hoyt, R. E. (2000). The demand for flood insurance: empirical evidence. *Journal of Risk and Uncertainty, 20*(3), 291-306.

Burgos, A. (2002). Learning to deal with risk: what does reinforcement learning tell us about risk attitudes? *Economics Bulletin, 4*(10), 1-13.

Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning*. Oxford, England: John Wiley & Sons.

Camerer, C., & Ho , T. H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica, 67*(4), 827-874.

Chiu, W. H. (2005). Skewness preference, risk aversion, and the precedence relations on stochastic changes. *Management Science, 51*(12), 1816-1828.

Cohen, M., & Sebstad, J. (2005). Reducing vulnerability: The demand for microinsurance. *Journal of International Development, 17*(3), 397-474.

Connor, R. A. (1996). More than risk reduction: The investment appeal of insurance. *Journal of Economic Psychology, 17*(1), 39-54.

Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review, 114*(1), 177-187.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science, 12*(5), 523-538.

Dionne, G., & Eeckhoudt, L. (1985). Self-insurance, self-protection and increased risk aversion. *Economics Letters, 17*(1–2), 39-42.

Ehrlich, I., & Becker, G. S. (1972). Market insurance, self-insurance, and self-protection. *Journal of Political Economy, 80*(4), 623-648.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112*(4), 912-931.

Ganderton, P. T., Brookshire, D. S., McKee, M., Stewart, S., & Thurston, H. (2000). Buying insurance for disaster-type risks: experimental evidence. *Journal of Risk and Uncertainty, 20*(3), 271-289.

Jaspersen, J. G., & Peter, R. (2014). The influence of adaptation on skewness seeking. *Munich Risk and Insurance Center Working Paper, No. 20*.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263-291.

Kriesel, W., & Landry, C. (2004). Participation in the National Flood Insurance Program: An empirical analysis for coastal properties. *The Journal of Risk and Insurance, 71*(3), 405-420.

Kunreuther, H. (1996). Mitigating disaster losses through insurance. *Journal of Risk and Uncertainty, 12*(2-3), 171-187.

Kunreuther, H., Onculer, A., & Slovic, P. (1998). Time insensitivity for protective investments. *Journal of Risk and Uncertainty, 16*(3), 279-299.

Laury, S., McInnes, M., & Swarthout, J. T. (2009). Insurance decisions for low-probability losses. *Journal of Risk and Uncertainty, 39*(1), 17-44. doi: 10.1007/s11166-009-9072-2

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*: John Wiley and Sons.

March, J. G. (1988). Variable risk preferences and adaptive aspirations. *Journal of Economic Behavior & Organization, 9*(1), 5-24.

March, J. G. (1996). Learning to be risk averse. *Psychological Review, 103*(2), 309-319.

March, J. G., & Shapira, Z. (1992). Variable risk preferences and the focus of attention. *Psychological Review, 99*(1), 172-183.

McClelland, G., Schulze, W., & Coursey, D. (1993). Insurance for low-probability hazards: A bimodal response to unlikely events. *Journal of Risk and Uncertainty, 7*(1), 95-116. doi: 10.1007/BF01065317

Meyer, R. J. (2012). Failing to learn from experience about catastrophes: The case of hurricane preparedness. *Journal of Risk and Uncertainty, 45*(1), 25-50.

Roth, J., McCord, M. J., & Liber, D. (2007). *The landscape of microinsurance in the world's 100 poorest countries*: MicroInsurance Centre.

Schade, C., Kunreuther, H., & Koellinger, P. (2012). Protecting against low-probability disasters: The role of worry. *Journal of Behavioral Decision Making, 25*(5), 534-543.

Shafran, A. P. (2011). Self-protection against repeated low probability risks. *Journal of Risk and Uncertainty, 42*(3), 263-285.

Slovic, P., Fischhoff, B., Lichtenstein, S., Corrigan, B., & Combs, B. (1977). Preference for insuring against probable small losses: Insurance implications. *The Journal of Risk and Insurance, 44*(2), 237-258.

Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning*: MIT Press.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*(4), 297-323.

Viscusi, W. K. (1979). Insurance and individual incentives in adaptive contexts. *Econometrica, 47*(5), 1195-1207.

Appendix

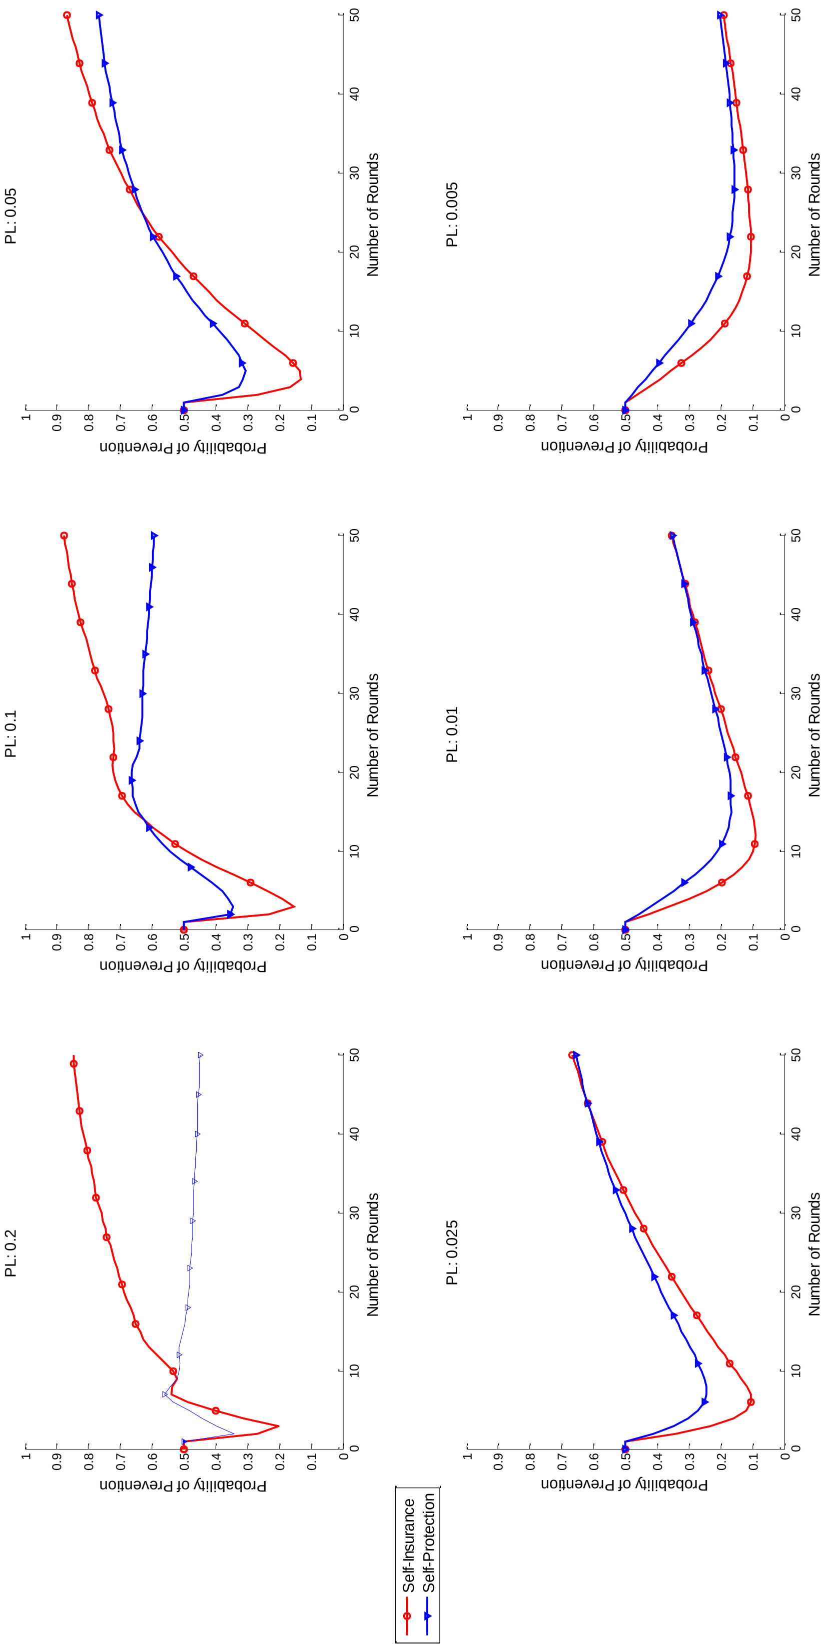

Figure 12: Combined prevention without loading: Probability to self-insure and to self-protect depending on time and probability of loss

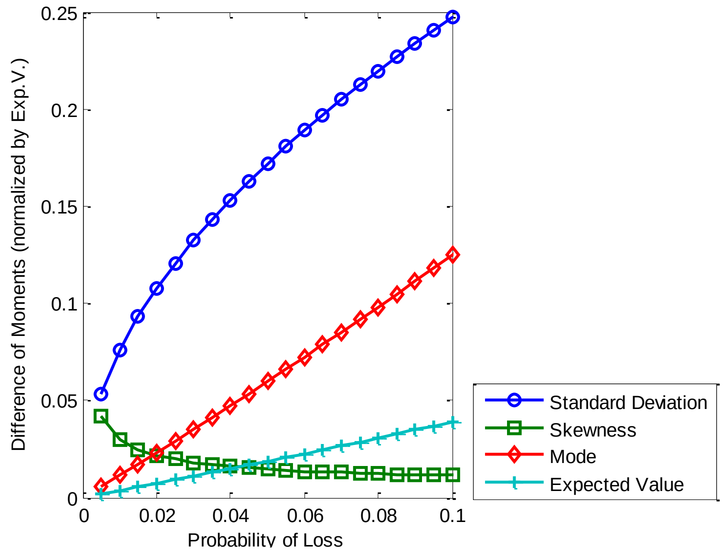

Figure 13: No Prevention and Combined Prevention: Difference of Distributional Properties

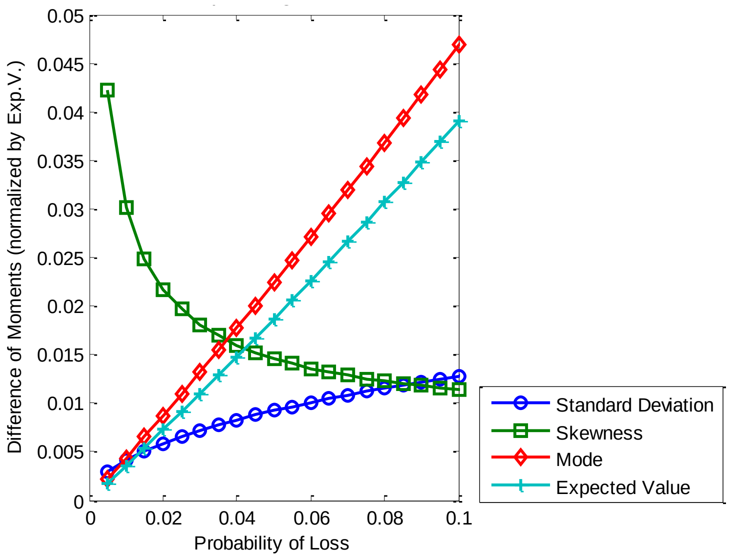

Figure 14: Self-Insurance and Combined Prevention: Difference of Distributional Properties

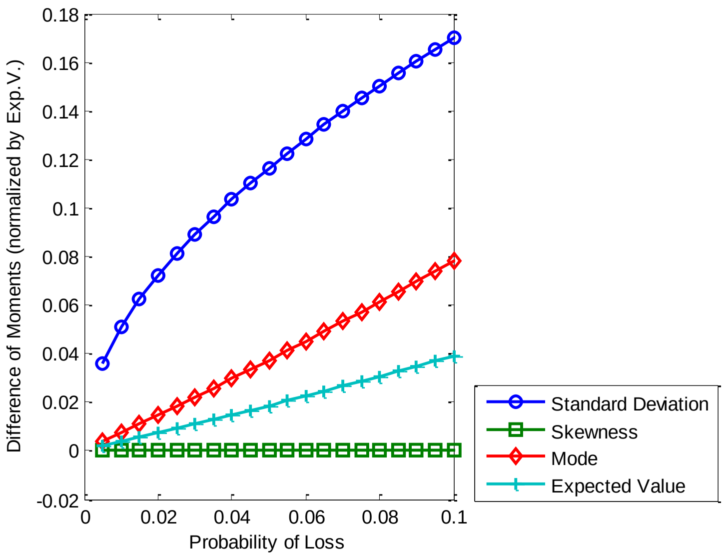

Figure 15: Self-Protection and Combined Prevention: Difference of Distributional Properties

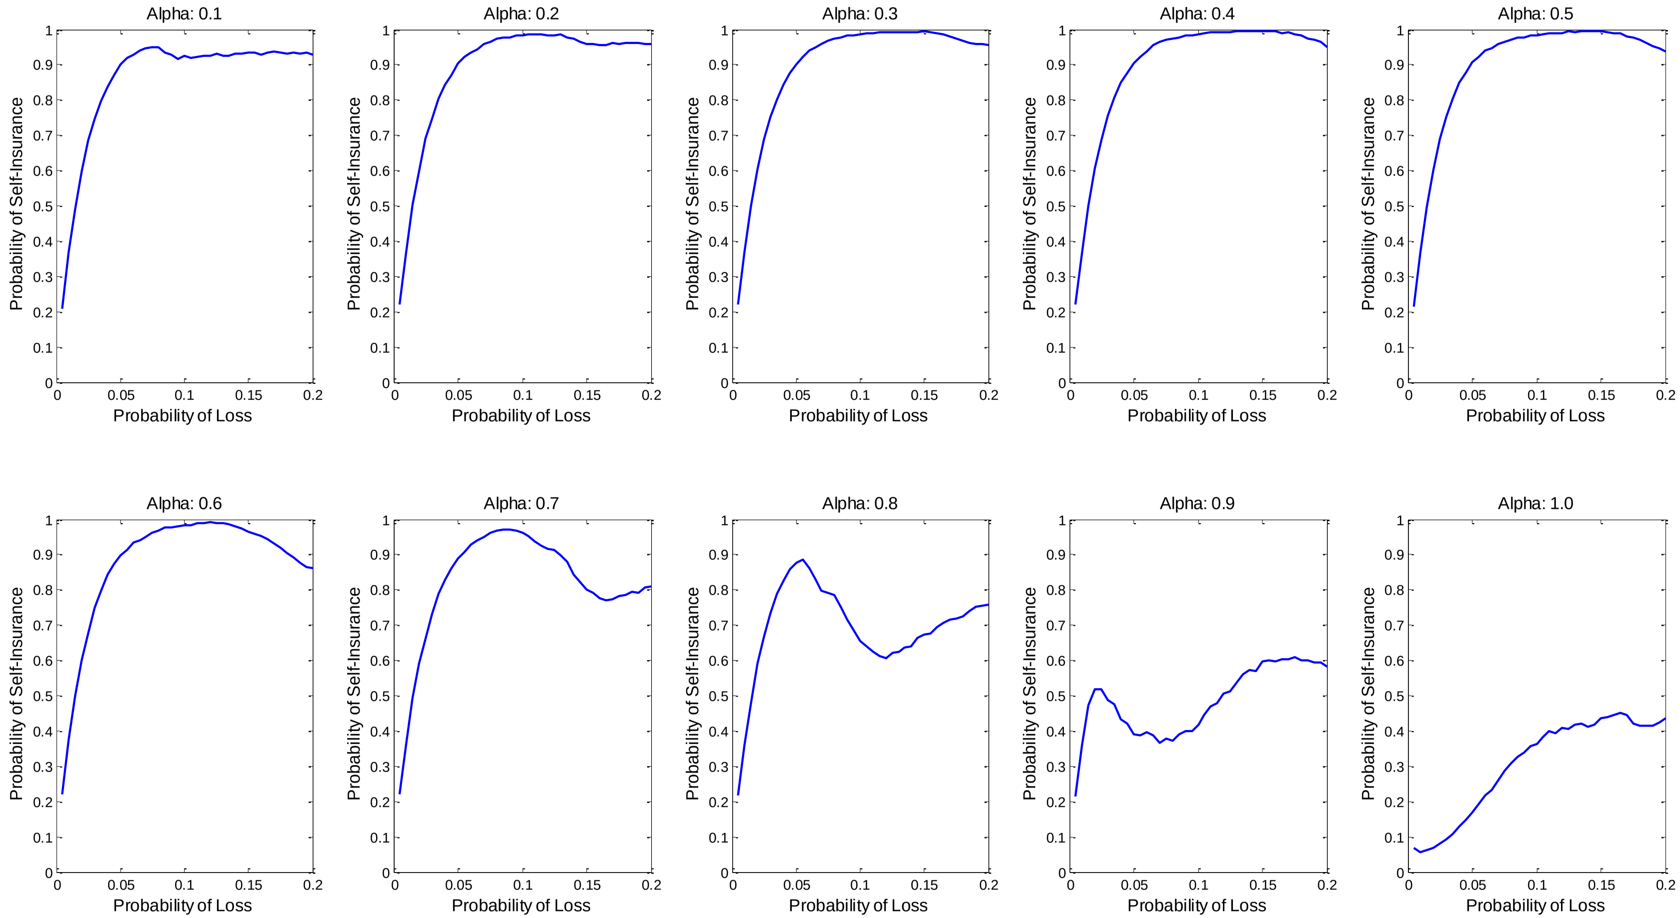

Figure 16: Probability to self-insure depending on reference outcome updating parameter alpha at beta = 0.6:

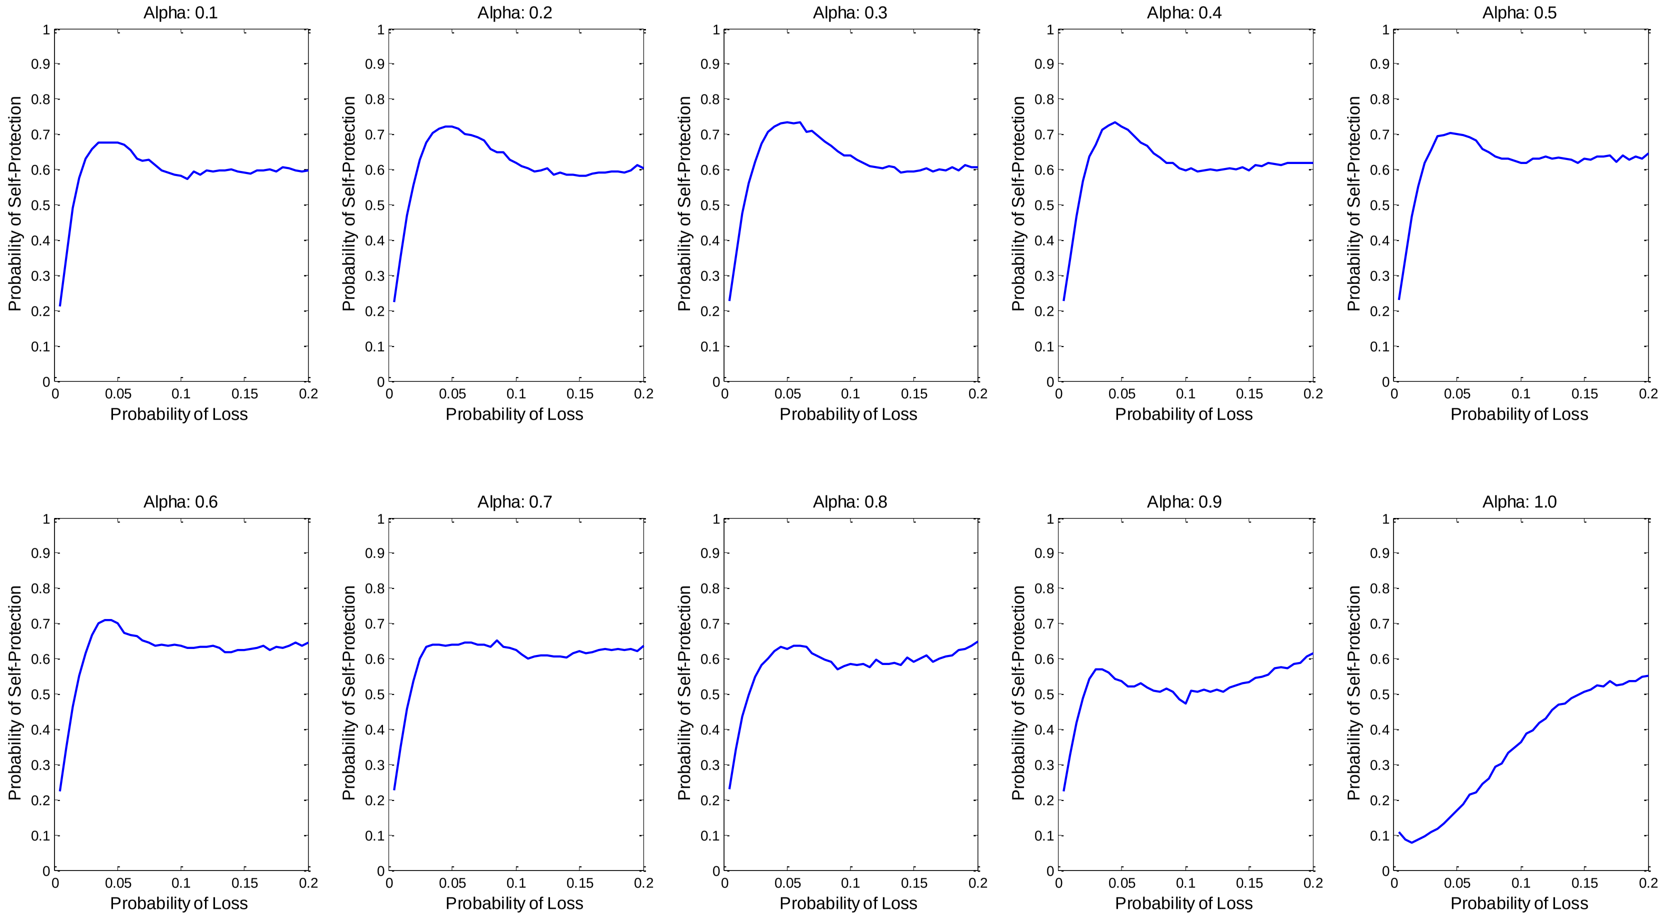

Figure 17: Probability to self-protect depending on reference outcome updating parameter alpha at beta = 0.6

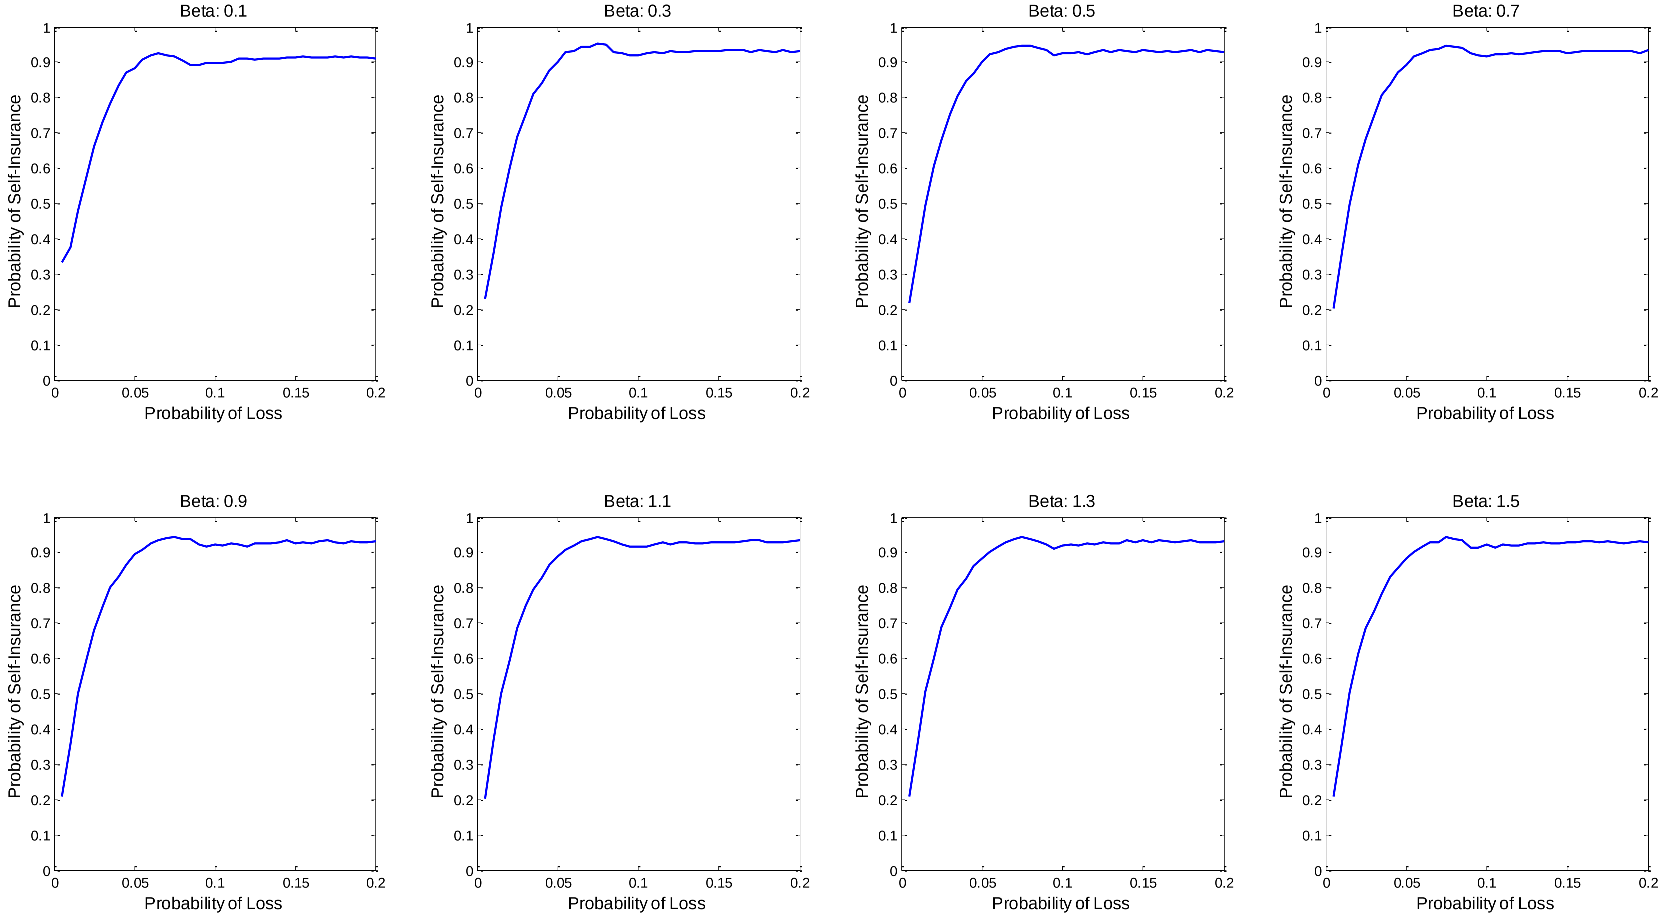

Figure 18: Probability to self-insure depending on learning parameter beta at alpha = 0.1

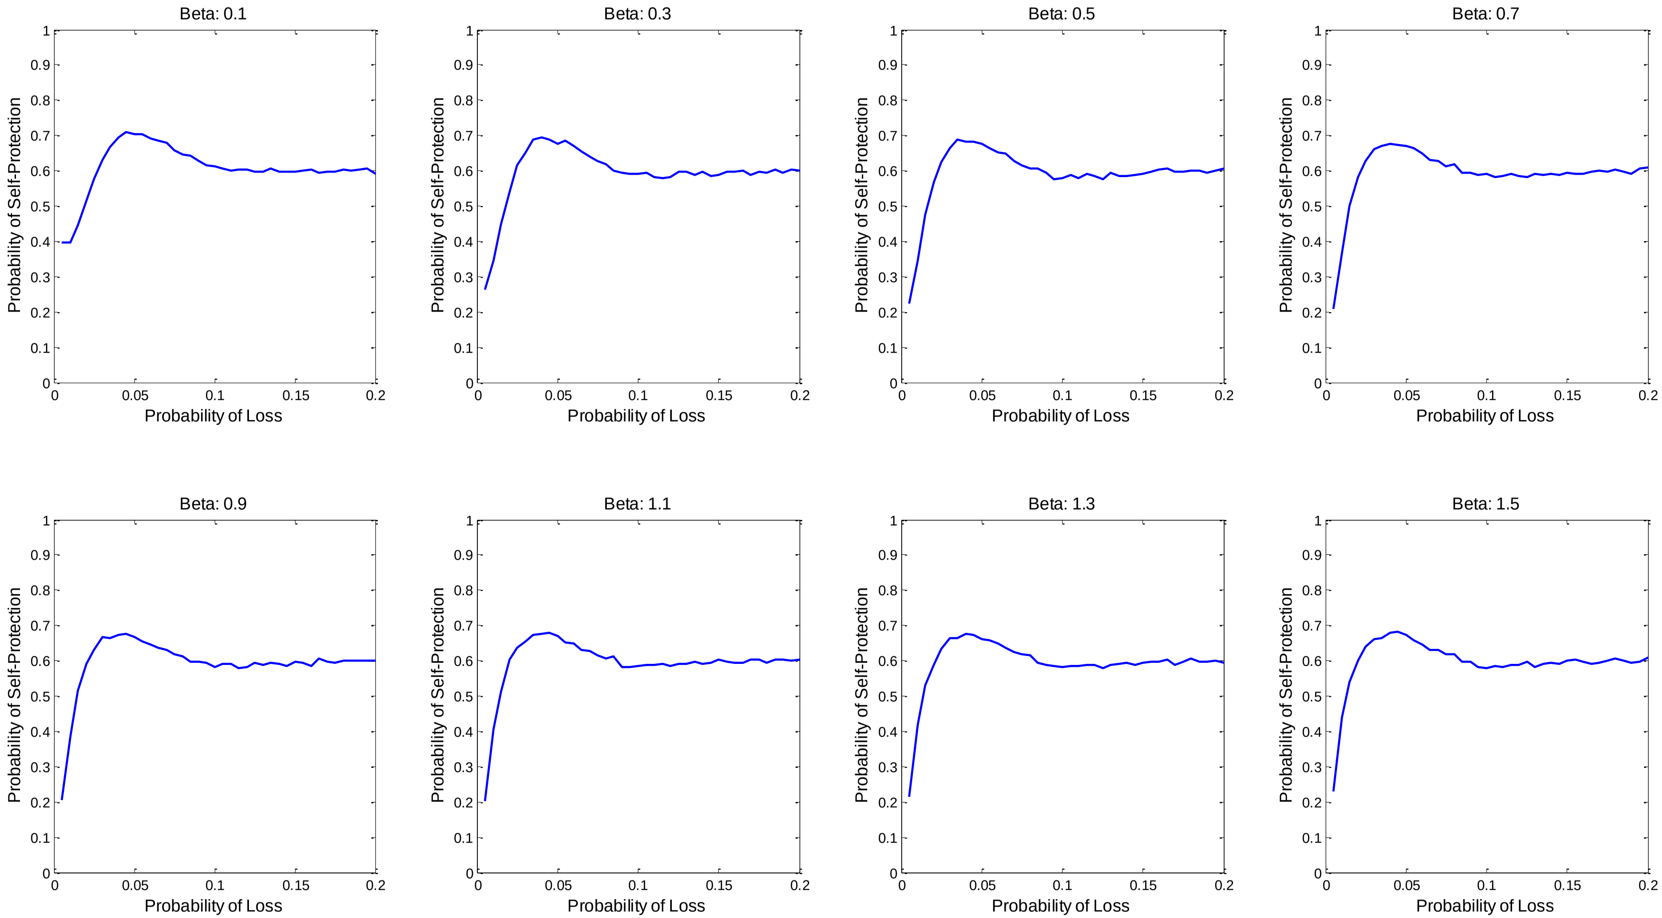

Figure 19: Probability to self-protect depending on learning parameter beta at alpha = 0.1

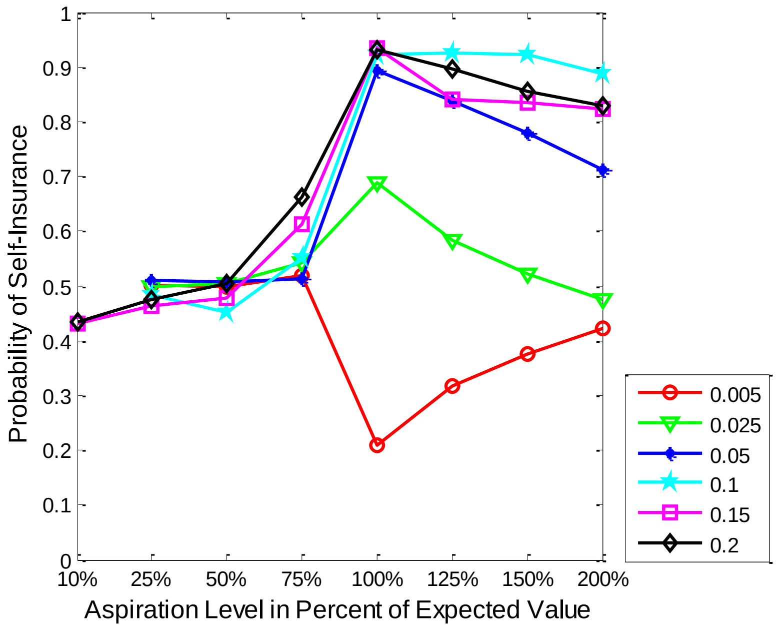

Figure 20 Probability to self-insure depending on initial aspiration level

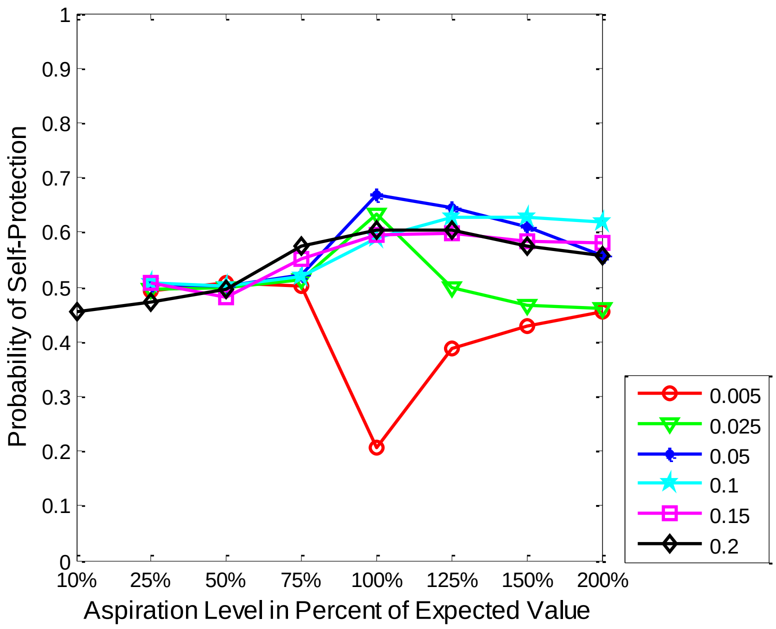

Figure 21 Probability to self-protect depending on initial aspiration level

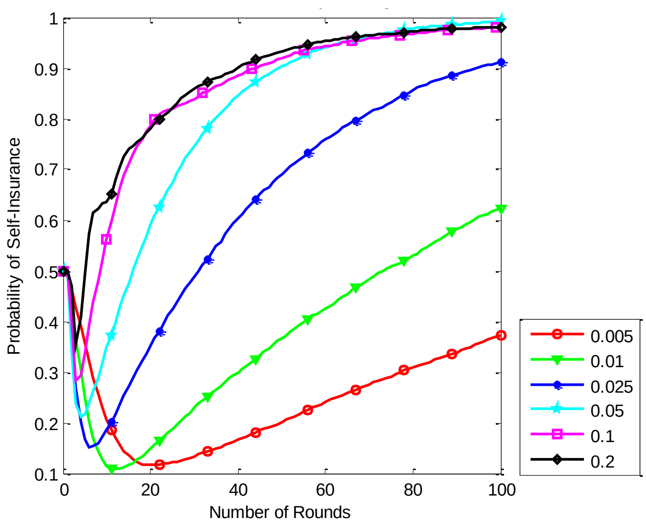

Figure 22 Probability to invest in self-insurance depending on probability of loss over 100 rounds

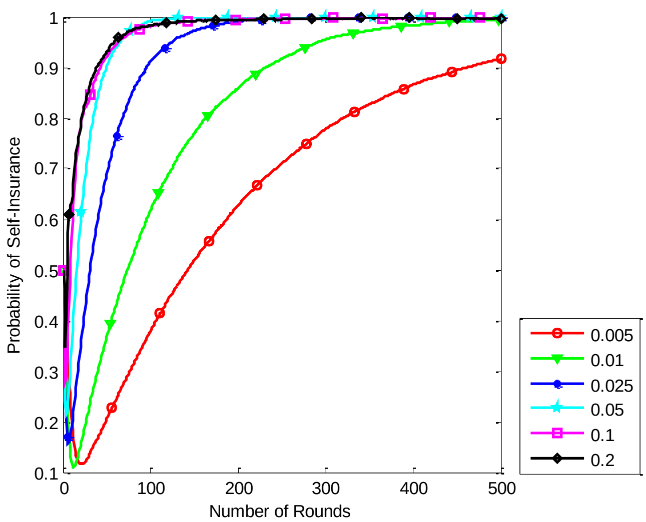

Figure 23 Probability to invest in self-insurance depending on probability of loss over 500 rounds

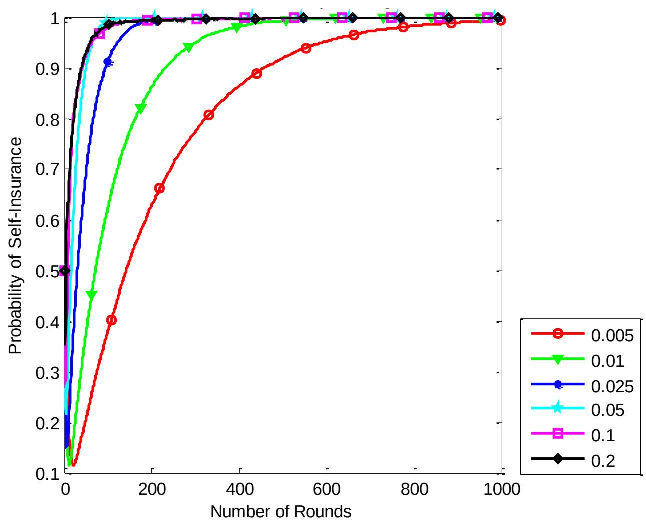

Figure 24 Probability to invest in self-insurance depending on probability of loss over 1000 rounds

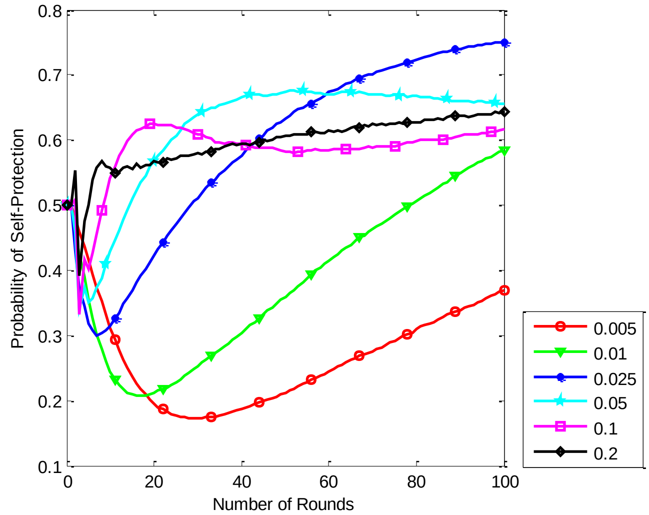

Figure 25 Probability to invest in self-protection depending on probability of loss over 100 rounds

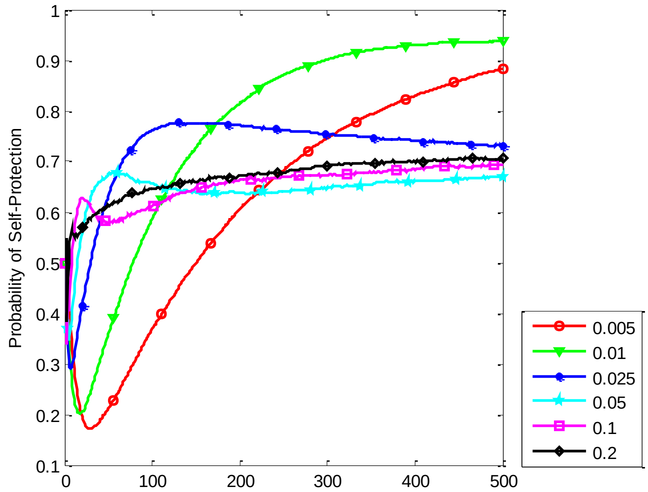

Figure 26 Probability to invest in self-protection depending on probability of loss over 500 rounds

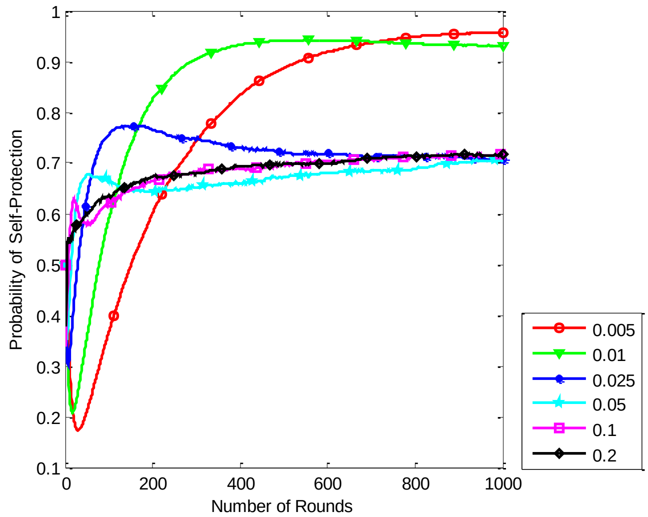

Figure 27 Probability to invest in self-protection depending on probability of loss over 1000 rounds

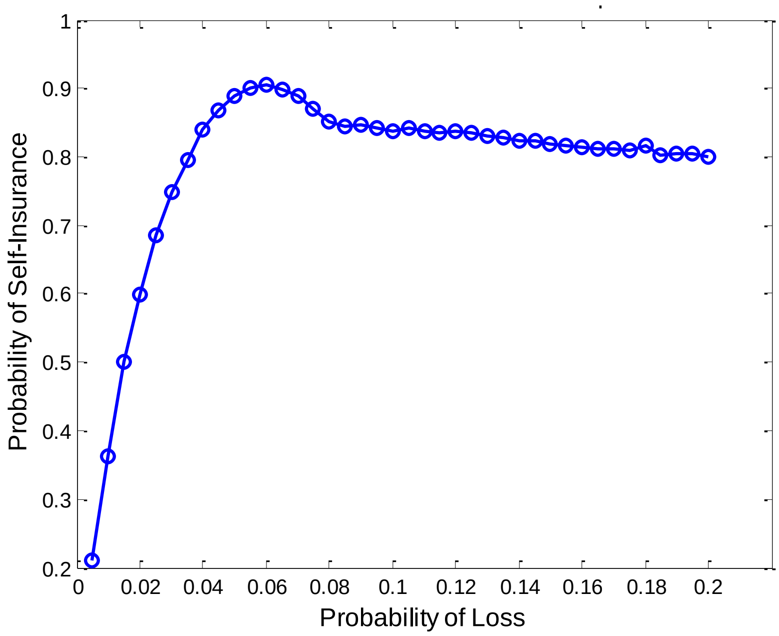

Figure 28: LPHC vs. HPLC: Probability to self-insure at t=50

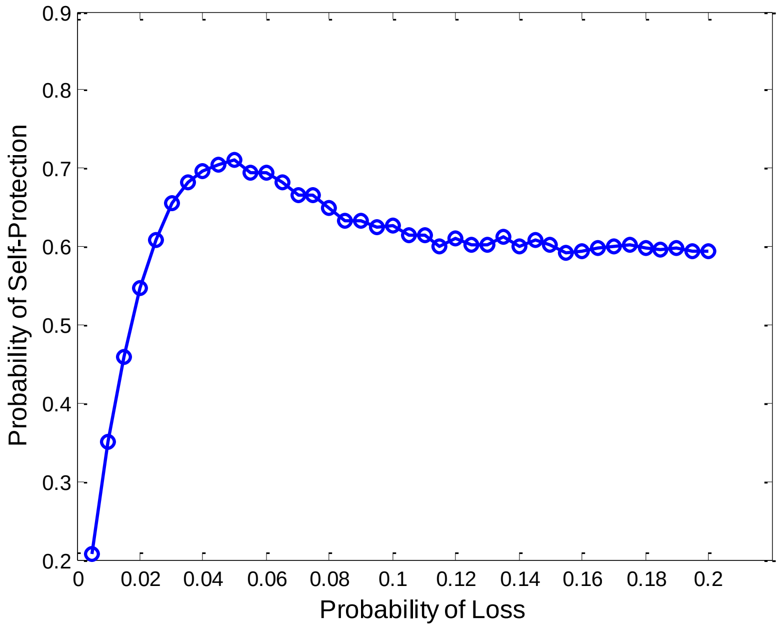

Figure 29: LPHC vs. HPLC: Probability to self-protect at t=50