Risky choices: a theory of preference over differences in probabilities

Thomas Barré
I4 - Index Insurance Innovation Initiative
University of California, Davis

Very preliminary and incomplete
Please do not cite or quote

October 25th, 2014

Abstract: This paper introduces a new framework for analyzing risk preferences that does not rely on Von-Neumann and Morgenstern's expected utility theorem. In particular, the proposed model does not require the (much debated) axiom of independence. In our setting, individuals choose between two lotteries L and M, depending on a cumulative probability gap: they compare the probability that the outcome W is lower than each possible value x for each lottery, and put a weight on each combination (x ; probability gap). Formally, the individual chooses between lotteries L and M depending on the sign of the cumulative weighted probability gap. It allows us to build a theory of decision under risk where risk aversion is not linked to the concavity of Bernoulli's utility function. We show that in this framework, it is possible to estimate risk preferences non-parametrically and provide some simulation illustrations.

Introduction

Since Von-Neumann and Morgenstern (VNM henceforth) seminal work in 1947, theories of decision under risk rely on the use of Bernoulli's utility function: individuals would exhibit some concave utility function of wealth. And Arrow-Pratt's measures of risk aversion are derived from this concavity. The identification of this utility function is often very challenging, so that a majority of studies prefer to impose some parametric restrictions, and risk preferences under the assumption that this parametric utility function is correct (up to an affine transformation). Constant relative risk aversion is particularly popular among applied researcher because it can be reduced to the estimation of a single parameter.

A second important characteristic of VNM's expected utility representation of preferences is the independence axiom which says that preferences are linear in probabilities. The axiom was criticized by

Allais (1953) and Savage (1954), and gave rise to rank dependent utility theory (Quigging, 1982) where individuals overweight extreme events.

The present paper offers a different approach to risk preferences that does not rely on Bernoulli's utility functions and extends the idea of preferences over extreme events introduced by Quigging (1982). In our setting, individuals choose between two lotteries L and M, by comparing some *cumulative weighted probability gap* over the entire interval of possible wealth levels. This approach allows us to provide a new graphical representation of risk preferences for any lotteries L and M, which brings up a simple non-parametric estimator of risk preferences. We provide some simulations results based this estimator.

The next section presents a short history of the theory of risk preferences and the experiments developed to reveal preferences and identify the several dimensions of risk preferences and beliefs. Then, we introduce our new framework and some interesting properties. Finally, we introduce our non-parametric estimator and some encouraging simulation results.

History of risk preferences theories and experimental measures

Expected utility theory

It has long been considered that when individuals are offered to choose between two gambles, they should pick the one with the highest expected value. In 1713, Nicolas Bernoulli first observed that such criterion could not describe actual individual decisions. He pointed to the fact that if someone was offered to play a game where the player wins 2^n dollars if a coin is flipped until a tails appear at stage n, the player would only be willing to pay a small price to enter the game, while the expected payoff is infinite. Daniel Bernoulli later offered an explanation of this paradox when he introduced the concept of *utility*. According to Daniel Bernoulli, the value of an item should not be based on its price, but on the utility it yields. In the case of money, a gain of a thousand dollars would be more significant to the poor than to the rich. Hence, individuals would exhibit some increasing and concave utility function of wealth. In the case of the logarithmic utility function propose by Daniel Bernoulli, a millionaire would be willing to pay up to $10.94 to play this game, while a person with $1,000 could only pay up to $5.94.

In 1947, Von-Neumann and Morgenstern introduced a set of four axioms supposed to characterize rationality. If an individual is rational, his preferences would satisfy these four axioms, he would have a utility function and he will always prefer actions that maximize expected utility. The four axioms of VNM-rationality are *completeness*, *transitivity*, *continuity*, and *independence*.

Axiom 1 (Completeness) assumes that an individual has well defined preferences: for any lotteries *L,M*, exactly one of the following holds: $L \prec M$, $M \prec L$, or $L = M$ (either M is preferred, L is preferred, or there is no preference).

Axiom 2 (Transitivity) assumes that preferences are consistent across any three options: If $L \preceq M$ and $M \preceq N$, then $L \preceq N$.

Axiom 3 (Continuity) assumes that there is a "tipping point" between being *better than* and *worse than* a given middle option: If $L \preceq M \preceq N$, then there exists a probability $p \in [0,1]$ such that $pL + (1-p)N = M$.

Axiom 4 (Independence) assumes that a preference holds independently of the possibility of another outcome: If $L \prec M$, then for any N and $p \in (0,1]$, $pL + (1-p)N \prec pM + (1-p)N$.

The theorem: For any VNM-rational agent (i.e. satisfying axioms 1–4), there exists a function *u* assigning to each outcome *A* a real number *u(A)* such that for any two lotteries, $L \prec M$ iff $E(u(L)) < E(u(M))$, where *E(u(L))* denotes the expected value of *u* in *L* (we'll abbreviate E(u(L)) to Eu(L)):

$$Eu(p_1 A_1 + \cdots + p_n A_n) = p_1 u(A_1) + \cdots + p_n u(A_n)$$

As such, *u* can be uniquely determined (up to adding a constant and a multiplying by a positive scalar) by preferences between *simple lotteries*, meaning those of the form $pA + (1 - p)B$ having only two outcomes. Conversely, any agent acting to maximize the expectation of a function *u* will obey axioms 1–4. Such a function is called the agent's **von Neumann–Morgenstern (VNM) utility**.

In the case of lotteries over monetary amounts, it appears that Bernoulli’s utility function is central to VNM’s theory. This centrality of Bernoulli’s work was further recognized in Arrow (1965) and Pratt (1964) who developed a **measure of absolute risk-aversion** (ARA) defined as

$$A(w) = -\frac{u''(w)}{u'(w)}$$

Rank Dependent utility

Soon after the publication of VNM’s work, several authors criticized their axioms. Using a simple experiment, Allais (1953) demonstrated that the axiom of independence was violated by a large share of the population when presented the following choices: In a first game, individuals were asked to choose between two gambles A1 and B1. In gamble A1, individuals have 100% chance to receive $1 million. In gamble B1, the have 89% change of winning $1 million, 10% chance to receive $5 million and 1% to get nothing. In a second game, these same individuals are asked to choose between lotteries A2 and B2. In lottery A2, they have 11% chance to receive $1 million, and 89% to receive nothing. In lottery B2, they have 10% chance of winning $5 million, and 90% to get nothing.

Experimental evidence shows that a majority of individuals would choose the lottery A1 in the first game, and lottery B2 in the second game. However, the combination of these two choices contradicts expected utility. Indeed, choosing lottery A1 in the first experiment implies that

$$U(\$1M) > 0.89U(\$1M) + 0.01U(\$0M) + 0.1U(\$5M)$$

While choosing lottery B2 in the second experiment implies

$$0.89U(\$0M) + 0.11U(\$1M) > 0.9U(\$0M) + 0.1U(\$5M)$$

which can be rearranged as

$$U(\$1M) < 0.89U(\$1M) + 0.01U(\$0M) + 0.1U(\$5M)$$

According to expected utility, rational individuals should not choose the lottery A1 together with the lottery B2. This result is known as Allais's paradox. And the source of this paradox is the axiom of independence.

The first solutions proposed to solve this paradox involved some subjective probabilities (Savage, 1954, Kahneman and Tversky 1979), individuals would overweight low probability events. But these approaches give rise to violations of first order stochastic dominance or to the axiom of transitivity. The solution to this issue was later introduced by Quigging (1982) who developed the idea of rank dependent utility functions where individuals overweight unlikely extreme outcomes. Transformations are applied to the cumulative probability distribution function, rather than to individual probabilities.

Formally, Rank Dependent Utility theory implies an increasing rearrangement of possible outcomes w such that $w_1 \leq w_2 \leq \cdots \leq w_n$, and the individual maximizes the objective function$V_i(\,)$

$$V_i(\boldsymbol{w}) = \int_{\mathbf{0}}^{\infty} \frac{\partial g_i\big(\Phi(x)\big)}{\partial x} u_i(x).\,dx$$

where $\Phi(x)$ is the probability that the outcome w is lower than some value x, $g_i(\,)$ is a probability weighting function such that $g_i(0) = 0$ and $g_i(1) = 1$.

Loss aversion

A second major critic to expected utility was recently formulated by Rabin (2000). Building on Arrow (1971)'s result that expected utility maximizers are arbitrarily close to risk neutral when stakes are arbitrarily small, Rabin shows that if an individual shows anything else than risk neutrality over moderate stakes, then expected utility theory implies unrealistic levels of risk aversion over large stakes. In particular, Rabin demonstrates that an individual would turns down a 50%-50% gamble of losing $100 and wining $110 would reject any 50%-50% gamble of losing $1000 and wining $X, whatever the value of X or the initial wealth of the individual.

Rabin's solution to this critic is to simply abandon Expected Utility Theory and adopt Kahneman and Tversky (1979, 1992)'s Prospect Theory. In their model, Kahneman and Tversky (1979) observe that possible losses weight more than potential gains when individual take risky decisions, and these asymmetries cannot be explained by income effects or decreasing risk aversion. Hence, individuals would not make decisions based on terminal wealth but on income with respect to a reference point. Loss aversion was later combined with rank dependent utility in Kahneman and Tversky (1992)'s Cumulative Prospect Theory to provide a better description of individual choices. Formally, this model writes:

$$V_i(\boldsymbol{w}) = \int_{-\infty}^{\mathbf{0}} \frac{\partial g_i^-\big(\Phi(x)\big)}{\partial x} u_i^-(x).dx + \int_{\mathbf{0}}^{+\infty} \frac{\partial g_i^+\big(\Phi(x)\big)}{\partial x} u_i^+(x).dx$$

Cumulative Prospect Theory is flexible enough to reconcile many of the departures from expected utility theory observed in individual choices. However, this flexibility is at the expense of simplicity; it is extremely difficult to identify $g_i^-(\,), u_i^-(\,), g_i^+(\,)$, and $u_i^+(\,)$ simultaneously. We come back to this question later in the paper.

Stochastic Dominance

A common characteristic that EUT and CPT share is that these theories assume individuals compare two gambles, by first giving a value to each gamble, and then compare these values together. Another approach (see for example Meyer (2014)) uses the concept of stochastic dominance. In this setting, a lottery of distribution *F* dominates another lottery of distribution *G* via

- *First order stochastic dominance* if: $F(x) \leq G(x)$ for all x. The implication for EUT is that $F(x)$ dominates $G(x)$ in the first degree if and only if $F(x)$ is preferred or indifferent to $G(x)$ by all EU decision makers with VNM utility functions satisfying $u'(x) \geq 0$ for all x.
- *Second order stochastic dominance* if: $\int_a^s \big(G(x) - F(x)\big)\, dx \geq 0$ for all $s \in [a, b]$. The implication for EUT is that $F(x)$ dominates $G(x)$ in the second degree if and only if $F(x)$ is preferred or indifferent to $G(x)$ by all EU decision makers with VNM utility functions satisfying $u'(x) \geq 0$ and $u''(x) \leq 0$for all x.

Recent studies (Eekhoudt et al.(2013), Deck and Schlesinger (2014)) demonstrate that individuals do exhibit preferences over low orders of stochastic dominance. While not directly related, the approach developed in this paper builds on this concept of stochastic dominance and comparisons of probabilities over the entire interval of possible values.

Experimental measurement

Beyond the description of observed behaviors, one objective of any theory of decision is to provide guidance about the kind of behavior individuals could exhibit in hypothetical conditions. In simple terms, we want to use these theories to predict individual behaviors if we changed the probability structure of their environment. For example, an insurer would like to know how many farmers would like to buy a new insurance product (and what coverage level) if this product was introduced in the market. We would like to rank individuals according to their level of risk aversion (and loss aversion, probability weighting, etc.) and know how each category of individuals will react to the introduction of this new insurance product.

There exist two main approaches to the measurement of risk preferences:

- Elicitation of the shape of individuals' utility functions, probability weighting functions and coefficients of loss aversion
- Elicitation of certainty equivalents

Here we only consider methods developed to elicit the shape of utility functions as it is more relevant to the topic of our own work. Recent years have seen the emergence of several very influential papers in this area (Biswanger (1981), Holt and Laury (2002), Eckel and Grossman (2002) or Tanaka et al. (2010)). Biswanger (1981), Holt and Laury (2002) and Eckel and Grossman (2002) propose to estimate a single coefficient of relative risk aversion for each individual, assuming that they behave according to Expected Utility Theory, while Tanaka et al. (2010) extend this approach to the estimation of a set of three parameters (constant relative risk aversion, loss aversion, and probability weighting) following insights from Cumulative Prospect theory. A growing literature employs these experimental methods to elicit risk preferences assuming EUT and CPT, and test which model best describes the data (Harrison and Rutström, 2009).

But all these methods assume some parametric forms of constant relative risk aversion, loss aversion and probability weighting, which might give misleading results. The reason why such parametric restrictions have been imposed lies in the way both theories have been developed: they both consider that individuals when individuals chose between two lotteries, they first evaluate each lottery, and then compare these values. Hence, if one wanted to write an econometric model on the probability than a VNM individual chooses a lottery L against a lottery M, he would have to estimate, assuming EUT framework:

$$Prob(L \geq M) = Prob\left(\int_{\mathbf{0}}^{\infty} \varphi_L(x).u_i(x).dx \geq \int_{\mathbf{0}}^{\infty} \varphi_M(x).u_i(x).dx\right)$$

$$= Prob\left(\int_{\mathbf{0}}^{\infty} \big(\varphi_M(x) - \varphi_L(x)\big).u_i(x).dx \leq 0\right)$$

Where $\varphi_j(x)$ is the probability to receive the outcome x in lottery j. Even this simple model, which ignores loss aversion and probability weighting, would be highly non-linear and extremely difficult to estimate non-parametrically. The theory presented below tries to solve this issue by building a framework that is (i) easy present graphically, and (ii) relatively easy to estimate and (iii) does not rely on the axiom of independence, nor on the use of Bernoulli's utility function.

A new theory of risk preferences: the cumulative weighted probability gap theory

Our theory of choice is inspired by the concept of stochastic dominance. We start from the simple observation that the worst possible lottery is to get a zero wealth for sure. And the best possible lottery, is characterized by a 100% chance to get an infinite wealth. Actually, if we compare two lotteries L and M each characterized by a certain amount (w_L and w_M), it appears that the concept of utility of wealth is unnecessary: $L > M$ if and only if $w_L > w_M$. This observation extends to any pair of lotteries L and M as long as one lottery is dominated by the other at the first order of stochastic dominance.

Now, if we compare two lotteries L and M where the expected value of lottery L equals the certain amount paid in lottery M, we know that an individual will prefer lottery M if and only if he is risk averse. Our objective now, is to draw this preference for certainty of higher payouts in a simple graph. To do so, we consider that individuals have preferences over probability gaps and wealth levels, and that preferences are not linear in (transformed) probabilities (one central assumption in EUT and CPT). We consider that individuals compare the probability to have a wealth level lower than w for each possible w, and they attribute a weight to each possible combination of a probability gap and wealth level.

Consider the case where an individual with initial wealth \$5,000. We offer him to play a 50%-50% chance game where he can either lose \$1000 or win \$2000. The expected gain in clearly positive, but the individual will agree to play the game only if the following condition is satisfied:

$$\int_0^{+\infty} h_i(\Phi_L(w) - \Phi_M(w); w)dw \leq 0$$

where $\Phi_L(w)$ is the cumulative density function of the lottery game, and $\Phi_M(w)$ is the cumulative density function if the individual does not play. $h_i(\,)$ is a function that gives a weight to each possible outcome w depending on the probability gap $\Phi_L(w) - \Phi_M(w)$.

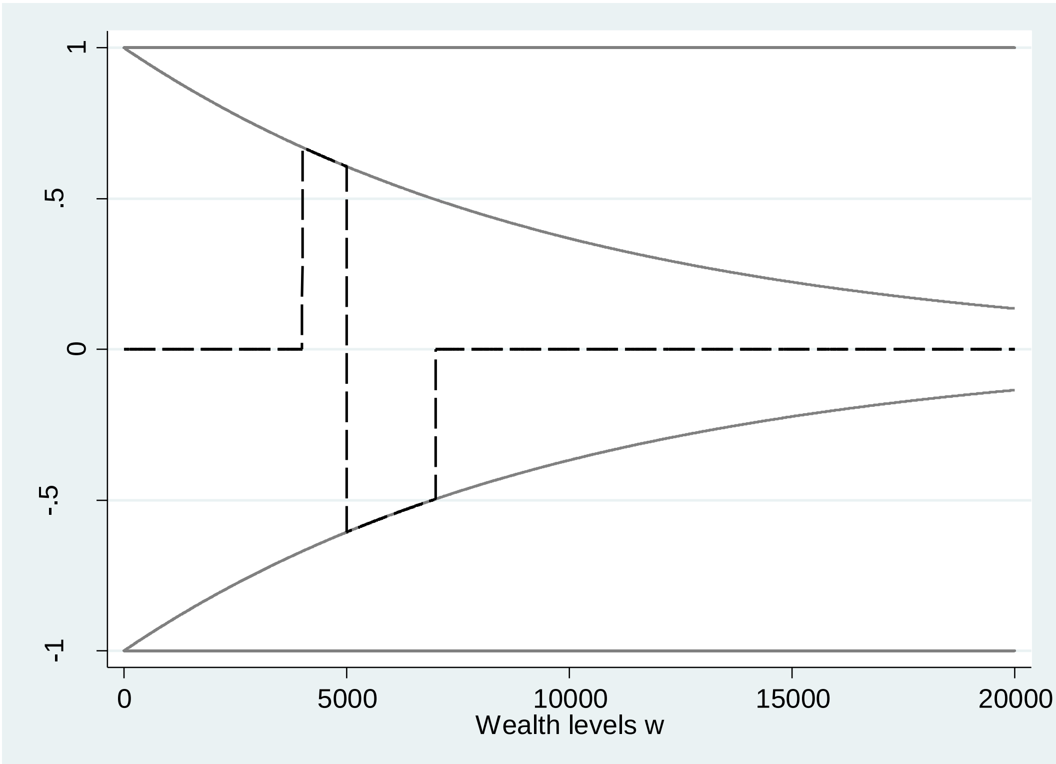

The solid grey curves represent the weights the individual gives to possible outcomes for which the probability gap is +50% (upper half) or -50% (lower half). The dashed line represents his preferences for the two lotteries.

In this case, we observe that the integral of the value function is negative. This means that the individual prefers to play the game; he values the potential gains more than the potential losses. If instead we offer the same individual to play a game where he has 60% chance to lose $1000 and 40% chance to win $2000, we obtain:

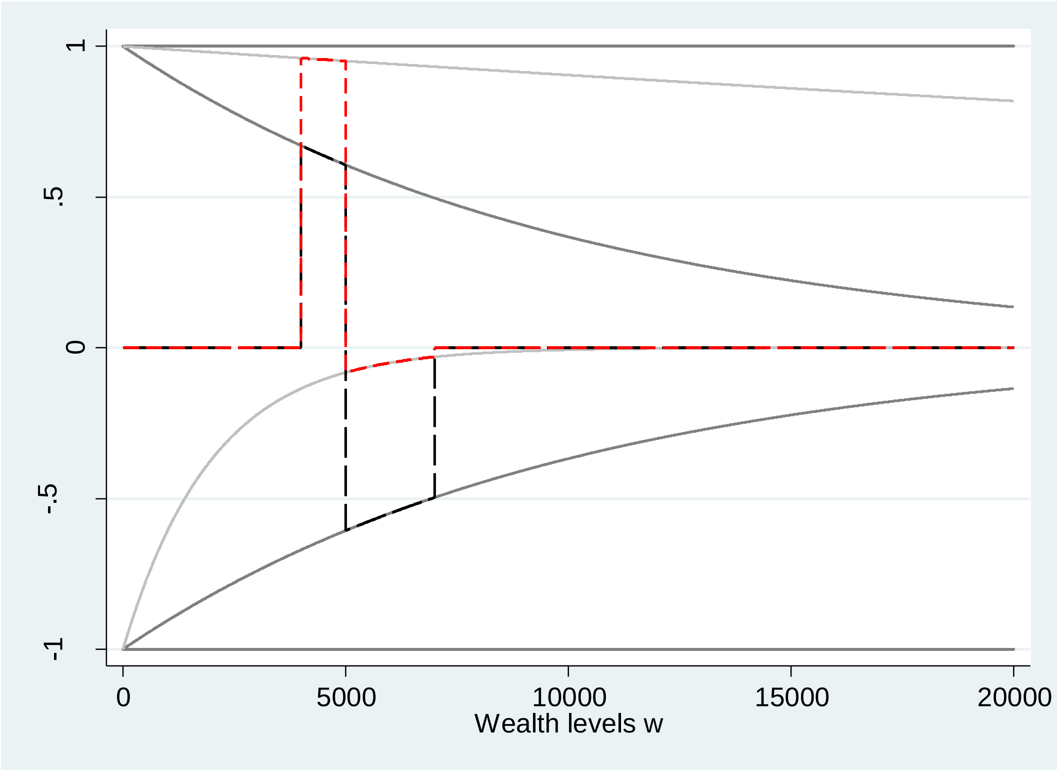

The solid dark grey curves represent the weights the individual gives to possible outcomes for which the probability gap is +50% (upper half) or -50% (lower half). The solid light grey curves represent the weights the individual gives to possible outcomes for which the probability gap is +60% (upper half) or -40% (lower half). The black dashed line represents preferences for the first 50%-50% game, and the red dashed line represents the preferences for the second 60%-40% game.

We observe that our player the integral for the second game is positive, so he would refuse to play game 2, even if the expected gain is still positive.

In our framework, for each possible outcome, an individual possesses an infinite number of "preference curves" whose shape depends on the probability gap for each outcome between the two lotteries considered. The individual does not compute an expectation of his utility (like in EUT or CPT), but instead chooses the dominant decision at the second order of stochastic dominance given his preferences.

A non-parametric estimator of risk preferences

This approach does not assume that individuals have an increasing and concave utility function of wealth, and it does not assume linearity in probabilities (i.e. the axiom of independence). Another nice characteristic of this approach is that we can estimate the preference curves for each probability gap non-parametrically. As mentioned before, the player chooses a lottery L over a lottery M if

$$\int_0^{+\infty} h_i(\Phi_L(w) - \Phi_M(w); w)dw \leq 0$$

Consider the example of game 1 above. We know that $\Phi_L(w) - \Phi_M(w) = 0$ for all w < 4000 & $w > 7000$, $\Phi_L(w) - \Phi_M(w) = 0.5$ for $4000 < w < 5000$ and $\Phi_L(w) - \Phi_M(w) = -0.5$ for $5000 < w < 7000$. The integral simplifies to:

$$\int_{4000}^{5000} h_i(0.5; w)dw + \int_{5000}^{7000} h_i(-0.5; w)dw \leq 0$$

If we assume that individuals do not exhibit loss aversion, the problem simplifies further to:

$$\int_{4000}^{5000} h_i(0.5; w)dw - \int_{5000}^{7000} h_i(0.5; w)dw \leq 0$$

If we apply Riemann's approximation method, we obtain:

$$Prob(L > M) = Prob\left(\sum_{j=1,\dots,n} (x_j - x_{j-1})h_i^{(50)}(t_j) - \sum_{k=1,\dots,m} (z_k - z_{k-1})h_i^{(50)}(t_k) + \varepsilon \leq 0\right)$$

where $4000 \leq x_{j-1} < t_j < x_j \leq 5000, x_0 = 4000, x_n = 5000,$, $5000 \leq z_{k-1} < t_k < z_k \leq 7000, z_0 = 5000, z_m = 7000$ and$h_i^{(50)}(x) = h_i(0.5; x)$

Assuming that Riemann's error terms are independent and identically distributed, and assuming a Gaussian distribution function, the problem becomes:

$$Prob(L > M) = \Phi\left(\frac{-\sum_j (x_j - x_{j-1})h_i^{(50)}(t_j) + \sum_k (x_k - x_{k-1})h_i^{(50)}(t_k)}{\sigma_\varepsilon}\right)$$

This is a probit model, where the index is a weighted sum of non-parametric functions. Hence, the model reduces to the estimation of a probit model with a semi-parametric index function. To illustrate this point, we again simplify our model assuming that we approximate each integral with only one point, in the middle of the interval, so that $n = 1$ and $m = 1$. We can write:

$$Prob(L > M) = \Phi\left(\frac{-(5000 - 1000)h_i^{(50)}(4500) + (7000 - 5000)h_i^{(50)}(6000)}{\sigma_\varepsilon}\right)$$

We can finally approximate $h_i^{(50)}$ by local polynomial regression. If each individual is offered to play several games similar to game1, involving 50%-50% gambles with different levels of gain and losses, we can estimate $h_i^{(50)}$. We simulate such case where an individual with initial wealth of $500 is offered to play 15 different games where losses are in a]$0.5;50[interval, and gains at least cover the losses (so that the expected gain is positive) and can go up to 2 times the potential loss.

We assume that our player has preferences of the form: $h_i^{(50)}(w) = (exp(-0.01w))$

The graph below shows the results of a very simple local polynomial estimator of degree 0 with a Gaussian kernel. While the results are not perfect yet, it is very encouraging. Indeed our simulation exercise is based on a very bad approximation of the integral, and our non-parametric estimator is simply not well suited for this kind of exercise where we have non-equidistant observations and the curve we want to estimate is convex. A second order polynomial approximation would perform much better.

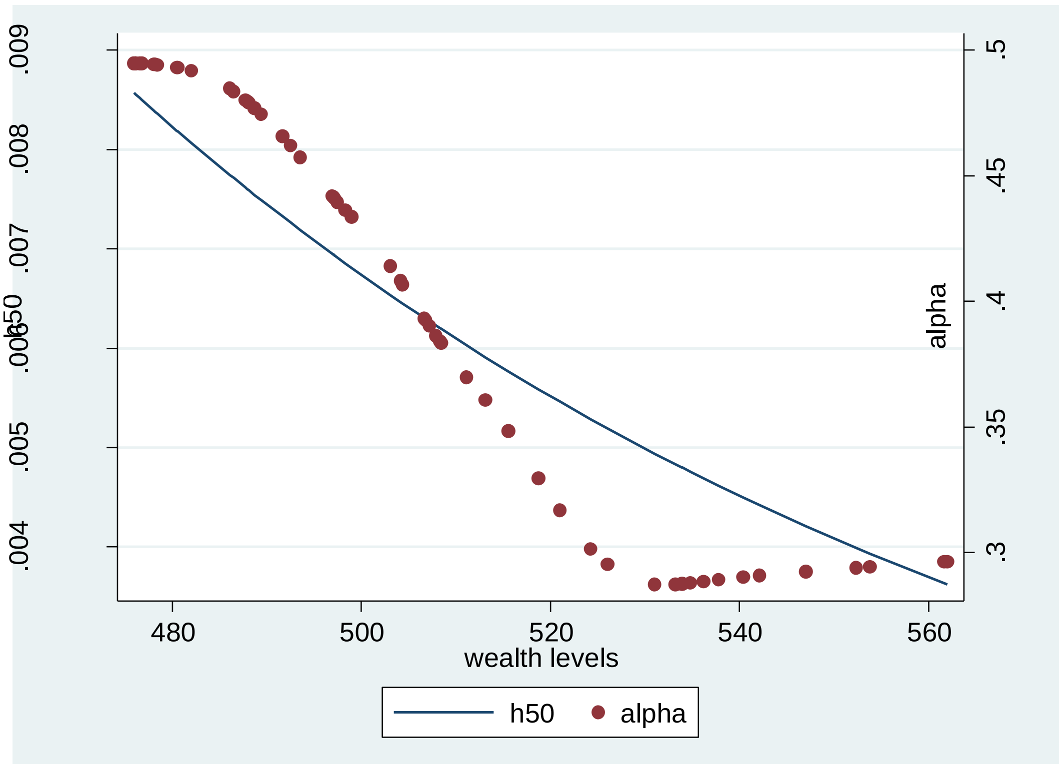

$h = h_i^{(50)}(w) = (exp(-0.01w))$; alpha = local polynomial estimator (degree 0)

Some additional theoretical results

….

Conclusion

….

References

….