UNIVERSITY *of York*

Discussion Papers in Economics

No. 15/03

Portfolio Choice Under Ambiguity

Enrica Carbone, Xueqi Dong and John Hey

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

Portfolio Choice Under Ambiguity

Enrica Carbone*, Xueqi Dong** and John Hey***

Abstract

This paper represents an intersection between two lines of research. The first is portfolio choice theory, which underlies much of finance; the second is the elicitation of preferences under uncertainty. The theory of the behaviour of financial markets builds heavily on portfolio choice theory; until recently this has assumed that preferences are of a particularly simple kind. In contrast research on preferences has revealed that people have more sophisticated preferences. This paper tries to bring the two fields together by investigating, in a portfolio choice context, the preferences that are revealed by decisions. In the second of these two fields, researchers are increasingly using allocation problems to elicit the preferences of subjects, believing that such problems are more informative, and perhaps more natural, than other elicitation methods. At the same time portfolio choice theory is itself concerned with an allocation problem. Usually in experimental finance the allocation problems are over Arrow securities each of which pays off only in one state of the world. Instead we study the more realistic case, familiar from finance, in which all assets pay off in all states of the world. To make our study more realistic we frame the problem as one under ambiguity, where the probabilities of the states are not known to the decision-maker. This enables us to compare the performance of some recent theories of behaviour under ambiguity as well as traditional ones (such as Mean-Variance) from the theory of finance. We also identify a 'rule of thumb' that decision-makers may be using in this rather complex scenario. This research may help us to understand more fully actual portfolio choice decisions.

* Seconda Universita degli Studi di Napoli (enrica.carbone@gmail.com)
** Royal Holloway University of London (xueqi.dong@rhul.ac.uk)
*** University of York (john.hey@york.ac.uk, corresponding author)

Keywords: ambiguity, portfolio choice, preferences under ambiguity, securities.
JEL codes: C9, D81, G02, G11

1. *Introduction*

The theory of financial markets builds heavily on portfolio choice theory. In turn this is constructed on assumptions about preferences under uncertainty. Until recently these assumptions have been that decision-makers have mean-variance preferences or that they are expected utility maximisers. However, in the research field of decision-making under uncertainty more sophisticated preference functionals have been proposed, and some appear to be empirically valid. This paper marries these two fields, by experimentally studying the portfolio choice decisions of subjects, and inferring from the data their underlying preferences. So we not only elicit preferences, we elicit them in a particularly relevant context.

There are several methods used by economists to elicit the preference functionals of subjects in situations of uncertainty. These include *Holt-Laury Price Lists* (Holt and Laury 2002), *Pairwise Choice questions* (Hey and Orme 1994) and the *Becker-DeGroot-Marschak (BDM) mechanism* (Becker *et al* 1964). More recently researchers have been using the *Allocation Method*, pioneered originally by Loomes (1991), revived by Andreoni and Miller (2002) in a social choice context, and later by Choi et al (2007) in a risky choice context. The allocation method seems potentially more informative than pairwise choice and price lists, and possibly easier to understand by subjects than the BDM mechanism.

In finance, allocation problems are familiar and are usually referred to as *portfolio choice* problems. However, in the language of finance experts, the kinds of portfolio choice problems presented to subjects in *experiments* are almost always over what are termed Arrow securities, *each* of which pays off some amount in *just one* state of the world. In practice, portfolio choice problems are over more realistic assets – *all* of which have some payoff in *each* state of the world. That is what we implement in our experimental investigation.

Clearly there is uncertainty (about which state of the world may occur) in practice. Most experimental studies characterise this uncertainty as *risk*, and subjects are told the probabilities of the various states occurring. In the real world, such probabilities are not known by the decision-makers so the uncertainty is better characterised as *ambiguity*. In our experiment we follow this route, as we shall describe.

Characterising uncertainty as ambiguity rather than risk opens up many possibilities for theorists, and the last decade has witnessed a proliferation of new theories of behaviour under ambiguity – most of which are surveyed in Etner *et al* (2012). We use our experimental data to examine a subset of these theories to see which best explains the behaviour of the subjects. At the same time theorists in finance cling to Mean-Variance analysis, despite it being inconsistent with much of decision theory. But Mean-Variance analysis is much simpler to apply than many modern theories of behaviour under ambiguity, and decision-makers may well use this rather than some more complicated method of taking decisions. With Mean-Variance analysis decision-makers have to guess at the underlying probabilities, assume that they are true, and then work out the mean and variance of any given portfolio to determine their optimal portfolio. Contrast this with what decision-makers are presumed to do under the modern theories – which we shall outline shortly. But what Mean-Variance postulates that decision-makers do may well be closer to what real decision-makers do: rather than complicate an already complicated problem with a complicated decision rule, they may simplify it by adopting a simple decision rule. Indeed one can push this argument further and argue that decision-makers adopt some 'rule of thumb'. *Ex ante*, of course, we may not know what such a rule may look like – but we can search for evidence of it in our data. So we test how well mean-variance explains behaviour, and we search for a 'rule of thumb'.

In summary, we present a set of portfolio choice problems under ambiguity to our subjects. We fit a variety of models to our data. We see which model of preferences best explains the data.

The paper is organised as follows. We start in section 2 by describing the kind of portfolio choice problems we presented to our subjects, and then look at various theories saying how people should behave. We start simple with a situation of risk and discuss the implications of Expected Utility (EU) theory and Mean-Variance theory. We then move on to ambiguity and look at MaxMin (Gilboa and Schmeidler 1989) preferences and then generalise them to α-MEU (Ghirardato *et al* 2004) preferences. In section 3 we anticipate the constraints we put on our subjects (one of which was that they were not allowed to hold negative amounts of any assets) and explore the implications, particularly from the point of view of deriving the optimal decisions under any particular preference functional. In section 4 we give more detail about our experimental design. In section 5 we discuss the appropriate stochastic specification underlying our econometric analysis, linking it to the constraints imposed on our subjects. Section 6 reports our results and section 7 concludes. Additional material can be found on the EXEC website.

2. The Portfolio Choice Problem and Possible Solutions

The decision-maker (DM) is given an endowment (which we normalise here to 100, as was the case in our experiment) in cash to allocate to three assets: one with a certain return (which we normalise to 1); and the other two with uncertain[1] returns, which depend upon which *state of nature* occurs. The number of such states is set at 3, which makes the problem a meaningful[2] one while reducing its complexity. Denote by c_1 and c_2 the allocations to the two uncertain assets 1 and 2. This implies that the allocation to the certain asset c_0 is given by $c_0 = 100 - c_1 - c_2$. Crucial to the allocation problem are the returns on the uncertain assets. Denoting by r_{ij} the *absolute* return on asset *i* if state *j* occurs, we have the following *returns table*:

	state 1	state 2	state 3
asset 1	r_{11}	r_{12}	r_{13}
asset 2	r_{21}	r_{22}	r_{23}

It follows that the portfolio payoff in state *j*, denoted by d_j, is given by $d_j = c_0 + r_{1j}c_1 + r_{2j}c_2$ $(j=1,2,3)$.

The DM's optimal allocations depend upon his or her preferences. If we start with Expected Utility (EU) theory under risk, or Subjective Expected Utility (SEU) under ambiguity, where p_j $(j=1,2,3)$ is the (subjective) probability of state *j* occurring, then the DM's objective function is the maximisation of $p_1u(d_1) + p_2u(d_2) + p_3u(d_3)$ where $u(.)$ is the individual's utility function. If instead the

[1] We are using this term at the moment to embrace both risk and ambiguity.

[2] If there were just 2 states there would not be enough information in the data to allow us to infer preferences.

DM follows Mean-Variance (MV) theory using probabilities p_j *(j=1,2,3),* then the objective is the maximisation of $\mu - r\sigma^2$, where r indicates the attitude to risk and the mean, μ, and variance, σ^2, of the portfolio are given by $\mu = p_1d_1 + p_2d_2 + p_3d_3$ and $\sigma^2 = p_1(d_1-\mu)^2 + p_2(d_2-\mu)^2 + p_3(d_3-\mu)^2$.

The above assumes that the DM works with either objective or subjective probabilities. If, however, the DM is in a situation of ambiguity and feels unable to attach unique probabilities to the various states of the world, then to model his or her behaviour we need to turn to one of the new theories of behaviour under ambiguity. In this paper we work with the simplest – MaxMin Expected Utility (MEU) and the α-MEU model. Both of these theories start by assuming that, while the DM cannot attach unique probabilities to the various states, he or she works with a *set of possible probabilities.* The theories do not say how this set is specified. We assume what appears to be the simplest: this set is all possible probabilities defined by (non-negative) *lower bounds* $\underline{p}_1$, $\underline{p}_2$ and $\underline{p}_3$ (where $\underline{p}_1 + \underline{p}_2 + \underline{p}_3 < 1$) on the probabilities. If you like, it is a little triangle properly within the Marschak-Machina triangle.

MEU postulates that the objective function of the DM is to choose the allocation which maximises the minimum expected utility over this set of possible probabilities. The α-MEU model generalises this to maximising the weighted average of the minimum and maximum expected utility over this set. More precisely, the α-MEU model's objective function is the maximisation of

$$\alpha\min(\underline{p}_1 \le p_1, \underline{p}_2 \le p_2, \underline{p}_3 \le p_3)[\, p_1u(d_1) + p_2u(d_2) + p_3u(d_3)] + (1-\alpha)\max(\underline{p}_1 \le p_1, \underline{p}_2 \le p_2, \underline{p}_3 \le p_3)[\, p_1u(d_1) + p_2u(d_2) + p_3u(d_3)].$$

MEU is the special case when $\alpha=1$.

Finally, we investigate a simple rule motivated in part by informally enquiring of the subjects how they had reached their decisions and in part by the data. We call this the Safety-First (SF) rule: allocations were made first such that their payoff in *all* states would be above some threshold $\underline{w}$ and then maximising the payoff in the most likely state[3]. When fitting this model, we estimate the parameter $\underline{w}$.

[3] It was clear from the Bingo Blower that there were more balls of one colour than either of the other two, though the precise numbers could not be known.

3. Constraints and Their Implications

In the experiment we did not allow the subjects to make negative allocations (which they might have wanted to do to maximise their objective function). We enforced this rule to stop the possibility of subjects making losses in the experiment. This meant that what we observe in the data are not optimal unconstrained allocations, but optimal *constrained* allocations. In order to fit the various models to the data we need to compute (for any given set of parameters) the optimal constrained allocations. While explicit analytical solutions are obtainable for the optimal unconstrained allocations for some of the preference functionals, they are not easily obtained for the optimal constrained ones. As a consequence we calculate them numerically.

There was also an additional 'constraint' on the allocations that subjects could make. In the experiment, the endowment in each problem was 100, and subjects were forced to implement allocations to the nearest integer. Given the non-negativity constraint this implied a set of 5151 possible allocations. Searching over these 5151 possible allocations proved to be a more efficient method of finding the optimal constrained allocations than using some built-in function, because of the complexity of the problem.

4. The Experimental Design

Subjects were presented with a total of 65[4] allocation problems, in each of which they were asked to allocate 100 experimental tokens to two possible assets or to keep some of them as tokens. In each of these they were shown a returns table. An example is the following:

	pink	green	blue
asset 1	1.7	0.9	0.6
asset 2	0	0.1	3.1

The colours relate to the way that ambiguity was implemented in the experiment. In the laboratory there was a Bingo Blower with pink, green and blue balls blowing around in continuous motion. Subjects could see the balls, and get a rough idea of the numbers and relative proportions of each

[4] These problems (and the number of them) were chosen after intensive pre-experimental simulations based on results from a pilot experiment, and were chosen to maximise the power of our estimates.

colour, so, when at the end of the experiment, one ball was ejected by them, they could not be sure of the probability of getting a ball of a particular colour[5]. Subjects were paid on a randomly chosen problem, with their payment being determined by the payoff (given their chosen allocations) for the state implied by the colour of a ball randomly ejected from the Blower.

A screen shot from the experiment can be seen in Figure 1; the 'returns table' was called the 'Payoff Table'. The triangle shows the set of all allowable allocations; as the subject moved his or her cursor around the triangle the Portfolio entries on the screen dynamically changed, and the implied payoffs for each colour were shown in the entries under 'Portfolio Payoff'. Subjects were forced to spend a minimum time of 30 seconds before registering their choice on any problem; there was a maximum time of 180 seconds per problem, and if they had not registered their choice by that time, it was taken to be an allocation of zero to the two uncertain assets. The instructions given to the subjects can be found here.

5. *Stochastic Specification*[6]

The object of the paper is to fit preference functionals to the experimental data and see which best explains the behaviour of the subjects. We do this subject by subject, as we believe that subjects are different. Our data are the actual allocations in each problem, denoted by x_1, x_2 and x_3 (where $x_1 + x_2 + x_3 = 100$)[7]. Each preference functional specifies, given the underlying behavioural parameters[8], an optimal constrained allocation on any problem. Let us denote these by x_1^*, x_2^* and x_3^*; again these add to 100. These depend upon the underlying behavioural parameters. It would be pleasing if $x_i = x_i^*$ for all i, for a particular preference functional and particular parameters, as this would enable us to identify the best preference functional. But this is unlikely to happen – the

[5] There were actually 10 pink, 20 green and 10 blue balls in the Blower, so the objective probabilities were 0.25, 0.5 and 0.25.

[6] This section can be safely skipped by those mainly interested in the results.

[7] We omit a subscript, indicating the problem number, for clarity

[8] These include the risk aversion index, the subjective probabilities, and, for the ambiguity models, the lower bounds on the probabilities. For SP the parameter is the threshold $\underline{w}$.

reason being, as is well-known, that subjects make errors when implementing their decisions[9]. So we need to admit the possibility of errors. We need also to model how these are generated. As both *x* and *x** are bounded (between 0 and 100) we proceed as follows. First we introduce new variables *y* and *y** which are the corresponding *x's* divided by 100. So $y_i = x_i/100$ and $y_i^* = x_i^*/100$ for $i=1,2,3$. These are bounded between 0 and 1. The obvious candidate distribution is the *beta* distribution which takes values over 0 and 1. Furthermore, it seems natural to begin with to assume that the actual allocations, whilst noisy, are not biased, so that each y_i has a mean of y_i^* (and hence that each x_i has mean x_i^*). Now a variable with a beta distribution has two parameters α and β, and the mean and variance of the variable are respectively $\alpha/(\alpha+\beta)$ and $\alpha\beta/[(\alpha+\beta)^2(\alpha+\beta+1)]$. Taking y_1 first, if we assume that its distribution is beta with parameters $\alpha_1 = y_1^*(s-1)$ and $\beta = (1-y_1^*)(s-1)$, this guarantees that the mean of y_1 is y_1^* and that its variance is $y_1^*(1-y_1^*)/s$. The parameter *s* here is an indicator of the *precision* of the distribution: the higher is *s* the more precise is the DM and the less noisy are the allocations.

Notice, however, that the variance of the distribution depends upon y_1^* – the closer it is to the bounds, the smaller it is, and at the bounds it becomes zero. This implies that this distribution cannot rationalise any non-zero allocation if the optimal is zero, nor can it rationalise any observation not equal to 1 if the optimal is 1. To get round this problem, we modify our definitions of the parameters α_1 and β_1. Instead of $\alpha_1 = y_1^*(s-1)$ and $\beta=(1-y_1^*)(s-1)$ we postulate that $\alpha_1 = y_1'(s-1)$ and $\beta = (1-y_1')(s-1)$ where $y_1'=b/2+(1-b)y_1^*$. There is a new parameter, *b*, which indicates the *bias* of the actual allocation, so that now the mean of y_1 is not y_1^* but instead $b/2+(1-b)y_1^*$. If *b* is zero then it is not biased, and as *b* increases the bias increases.

Now we turn to y_2. We must take into account that this must be between 0 and $1-y_1$. Hence $y_2/(1-y_1)$ is between 0 and 1. Here again a beta distribution is the natural candidate and we assume that the distribution of $y_2/(1-y_1)$ is beta with parameters α_2 and β_2 given by $\alpha_2 = y_2'(s-1)/(1-y_1)$ and $\beta_2 = (1-y_2')(s-1)/(1-y_1)$ where $y_2'=b/2+(1-b)y_2^*$. Clearly if $y_1 = 1$, this method is not applicable, and so in

[9] An alternative is that none of the preference functionals explain behaviour.

this case we assume that the error is made solely on y_1. In all cases the third allocation, y_3, is the residual.

Finally, in order to proceed to the likelihood function we should remember that allocations could only be made in integers. We assume that subjects round their intended allocations. So, for example, the likelihood of an observation of x_1 is equal to the cumulative probability from x_1-0.5 to x_1+0.5. The general form of the sum of log-likelihood function for all 65 problems can therefore be written as

$$\mathcal{L} = \sum_{j}^{65} log(L_1 L_2)$$

Here

$$\begin{cases} L_1 = F\left(\frac{x_1 + 0.5}{100}, \alpha_1, \beta_1\right) - F\left(\frac{x_1 - 0.5}{100}, \alpha_1, \beta_1\right) \\ L_2 = F\left(\frac{x_2 + 0.5}{100 - x_1}, \alpha_2, \beta_2\right) - F\left(\frac{x_2 - 0.5}{100 - x_1}, \alpha_2, \beta_2\right) \ when \ x_1 \neq 100 \\ 1 \ when \ x_1 = 100 \end{cases}$$

where $F(x,\alpha,\beta)$ is the *cdf* of a beta distribution with parameters α and β. These parameters are specified above.

We use the Maximum Likelihood Routine in Matlab to find the optimal estimates for our parameters (which are *r*, *s, b* the underlying probabilities or the lower bounds on them), and the goodness-of-fit of the various preference functionals.

6. Results

We have explored a number of different specifications and we report here just the best. Our primary concern is about the best fitting preference functional; we start with that. We measure the goodness-of-fit by the maximised log-likelihood, but we need to correct for the number of parameters in the preference functional – the number of degrees of freedom in the estimation.

We have already talked about the preference functionals we have fitted. Each of these involves a utility function; we have taken[10] this to be the *constant absolute risk aversion* form so that utility

[10] We also tried the CRRA utility function but this generally fitted worse.

u(x) is proportional to $-e^{-rx}$. In order to compare the goodness-of-fit of the different specifications, we need to distinguish between pairs of preference functionals one of which is nested within the other, and pairs of preference functionals where neither is nested within the other. We use the *Likelihood Ratio Test (LRT)* for the former and the Clarke test for the latter. We note that EU is nested within both MEU and α-MEU and that MEU is nested within α-MEU, but that none of the other functionals are nested within any other.

When one model is nested within another, the test statistic is

$$\mathrm{T} = 2(\mathcal{L}_1 - \mathcal{L}_0)$$

where $\mathcal{L}_0$ is the maximised log-likelihood of the nested model and $\mathcal{L}_1$ is the maximised log-likelihood of the nesting model. The test statistic has a Chi-square distribution with degrees of freedom equal to the difference in the number of parameters in the two competing models. As α-MEU has one more parameter than MEU and as MEU has one more parameter than EU, the corresponding degrees of freedom for EU *v* MEU, EU *v* α-MEU and MEU *v* α-MEU are 1, 2 and 1. The results are summarised in Table 1, which reports the percentage of the subjects[11] for which the test was significant.

Table 1: Percentage of subjects significant using the Likelihood Ratio Test

	significant at 5%	significant at 1%
MEU *v* EU	17%	13%
α-MEU *v* MEU	11%	7%
α-MEU *v* EU	21%	12%

Both MEU and α-MEU do moderately better than EU for a small number of subjects, which may not be surprising as the decision problem was one under ambiguity rather than under risk. Nevertheless EU performs well.

When models are *not* nested one within the other we use the *Clarke Test* (Clarke 2007). The null hypothesis is that the models are equally good, and hence on a particular problem the

[11] These are over 75 subjects. We had a total of 77 but we omit 2 from this analysis as they were extremely risk-averse, investing nothing in either risky asset; clearly all the models, with appropriate parameters, can equally well describe the behaviour of these subjects.

probability of the log-likelihood for one model being larger than the probability of the other model is ½. That is:

$$H_0:\ P(L_1 - L_2 > 0) = 0.5$$

Here L_1 and L_2 are the individual log-likelihoods of the 65 problems, which are calculated using the estimated parameters of the two competing models. The test statistic is

$$T = \sum_{i}^{65} I_i(L_1 - L_2)$$

Here

$$I_i = \begin{cases} 1, & L_1 - L_2 > 0 \\ 0, & L_1 - L_2 \leq 0 \end{cases} .$$

Under the null hypothesis *T* has a binomial distribution with parameters *n=65* and *p=0.5.* Thus an observation greater than 40 or less than 25 rejects the null hypothesis at the 5% significance level. The results are summarised in Table 2. These are the *percentages* of the 75 subjects for whom the test was significant.

Table 2: Clarke Test

(A) Comparisons between SF, EU, MEU and α-MEU (5% significance level)

EU *v* SF			MEU *v* SF			α-MEU *v* SF		
EU better than SF	SF better than EU	Neither better than the other	MEU better than SF	SF better than MEU	Neither better than the other	α-MEU better than SF	SF better than α-MEU	Neither better than the other
64%	9%	25%	64%	5%	31%	65%	5%	30%

(B) Comparisons between MV, and EU, MEU and α-MEU (5% significance level)

EU *v* MV			MEU *v* MV			α-MEU *v* MV		
EU better than MV	MV better than EU	Neither better than the other	MEU better than MV	MV better than MEU	Neither better than the other	α-MEU better than MV	MV better than α-MEU	Neither better than the other
48%	7%	45%	48%	4%	48%	52%	5%	43%

From the above, it seems that SF does not do particularly well and perhaps can be dismissed. More importantly, since much of finance theory is built on it, MV seems to be outperformed by EU and its generalisations for many of the subjects.

As a side issue, it may be interesting to report on the estimated probabilities for EU and the estimated lower bounds on the probabilities for MEU and α-MEU; recall that the true probabilities were 0.25 (pink), 0.5 (green) and 0.25 (blue). The averages (over all subjects) of the estimated probabilities for EU were 0.262, 0.530 and 0.208, which are very close to the true probabilities (though there was considerable dispersion across subjects). For MEU the average lower bounds were 0.228, 0.507 and 0.190, while for α-MEU they were 0.212, 0.490 and 0.171. These are (necessarily) lower than the corresponding EU probabilities, but only marginally so. These figures suggest that while, for some subjects, MEU or α-MEU are *statistically* superior to EU, the economic importance is marginal.

While SF does not perform particularly well, it may be if interest to report the estimated values of the threshold $\underline{w}$ – the distribution is in Figure 2. It will be seen from this that many subjects had a very high threshold – some approaching 100%. This alternatively could be interpreted as the result of very high risk-aversion, but this will of course by picked up by EU (or MEU or α-MEU) with a high estimated level of risk-aversion.

7. Conclusions

The main conclusion from the experiment is that MV did rather badly as an explanation of behaviour; this is rather worrying for finance theory. In contrast EU does rather well, not only compared to MV, but also compared with the generalisations, MEU and α-MEU: for relatively few subjects do these latter perform better. This indicates that subjects do not use a more complicated preference functional when choosing their allocations in a complicated setting. At the same time our 'simple' rule, SF, does worse than EU, suggesting some sophistication in subjects' decisions.

References

Andreoni, J. and Miller J, (2002), “Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism”, *Econometrica*, 70, 737-753.

Becker G.M., DeGroot M.H and Marschak J. (1964), “Measuring Utility by a Single-Response Sequential Method”, *Behavioral Science*, 9, 226-231.

Choi S., Fisman R., Gale D. and Kariv S. (2007), “Consistency and Heterogeneity of Individual Behavior under Uncertainty”, *American Economic Review*, 97, 1921-1938.

Etner, J., Jeleva, M. and Tallon J.M. (2012), “Decision Theory under Ambiguity”, *Journal of Economic Surveys*, 26, 234-270.

Ghirardato P., Maccheroni F. and Marinacci M. (2004), “Differentiating Ambiguity and Ambiguity Attitude”, *Journal of Economic Theory*, 118, 133–173.

Gilboa I. and Schmeidler D. (1989), “MaxMin Expected Utility with a non-Unique Prior”, *Journal of Mathematical Economics*, 18, 141-153.

Hey J.D. and Orme C. (1994), “Investigating Generalisations of Expected Utility Theory Using Experimental Data”, *Econometrica*, 62, 1291-1326.

Holt C.A. and Laury S.K. (2002), “Risk Aversion and Incentive Effects”, *American Economic Review*, 92, 1644-1655.

Loomes G. (1991), “Evidence of a new violation of the independence axiom”, *Journal of Risk and Uncertainty*, 4, 91-108.

Figures

Figure 1: A screen shot from the experiment

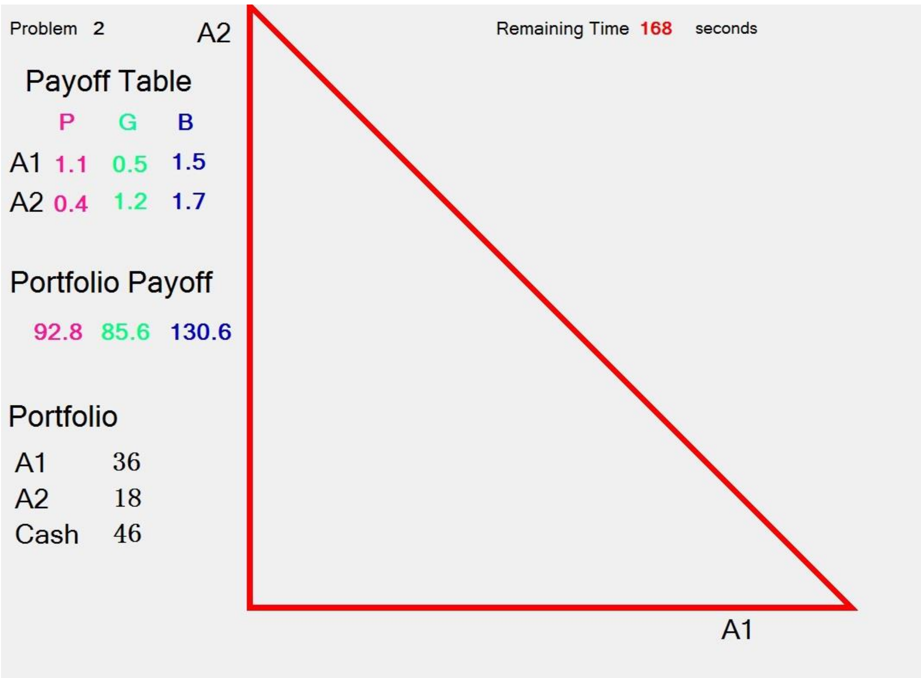

Figure 2: The distribution of the estimated threshold <u>*w*</u> for SF