## Title

Decision-makers minimize regret when calculating regret is easy

## Permalink

https://escholarship.org/uc/item/51t5×280

## Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

## Author

Srivastava, Nisheeth
Publication Date
2019

Peer reviewed

# Decision-makers minimize regret when calculating regret is easy 

Nisheeth Srivastava<br>Department of Computer Science, IIT Kanpur<br>India


#### Abstract

This paper provides empirical evidence that human decisionmakers use prospective regret minimization as their dominant decision strategy when regret calculations are cognitively easier to perform, and use expected utility maximization when they aren't. We designed a simple decision problem wherein utility maximization and expected regret minimization yield distinctly difference choices, and manipulated the cognitive effort involved in making regret calculations across respondent samples to arrive at our results. While previous research has associated ecological considerations like sense of responsibility and familiarity with this difference, we show that, at least in experimental settings, cognitive calculability in regret space appears to predominantly drive this difference. We also show that this preference for regret minimization can be countermanded by changing the distribution of options presented to the respondent, posing a challenge to simple sequential accounts of strategy selection learning which sequence strategy selection and application in order.


Keywords: decision-making; cognitive heuristics; cognitive effort; regret minimization; utility maximization

## Introduction

Regret is an important variable in humans' decision-making. Empirical investigations spanning psychology (Zeelenberg 1999; Connolly \& Reb, 2005), neuroscience (Coricelli et al., 2005) and economics (Loomes \& Sugden, 1982; Sarver, 2008) have demonstrated that in several decision contexts, humans behave as if they are trying to minimize prospective regret, rather than minimize prospective expected utility.

This distinction is of great significance for choice models that wish to track consequential human decisions. For instance, Chorus and colleagues have published a series of papers showing that a discrete choice model designed assuming regret minimization as the underlying choice strategy outperforms conventional random utility models (RUM) style discrete choice models in predicting future travel demand (Thiene, Boeri \& Chorus, 2012).

At the same time, conventional RUM models, assuming implicit utility maximization have proved their value in modeling human choices in a large array of applications (Small \& Rosen, 1981), suggesting that utility maximization is a useful approximation for peoples' intentions in such situations. Consequently, it is important to attempt to characterize situations wherein decision-makers are likely to prefer either of these decision-making strategies. Zeelenberg \& Pieters (2007) have suggested, on theoretical grounds, that regret-minimization is more likely to be used:
(a) when choices are perceived to be important and difficult,
(b) when the decision-maker expects to be held accountable for their choice and
(c) when the decision-maker anticipates receiving feedback about options in the near future.

There is also some empirical evidence supporting the basic premise that domain unfamiliarity may drive the use of regret minimization strategies, a mechanism that is substantially congruent with the theoretical factors identified by Zeelenberg \& Pieters (2007). Boeri, Scarpa \& Chorus (2014) have showed using discrete choice modeling on a transport choice dataset that the behavior of respondents unfamiliar with the choice context was better explained by regret minimizing models.

A common thread between such theoretical and empirical observations is the notion that regret is explicitly calculated by the respondent (Zeelenberg \& Pieters, 2007). It is because of this commitment to explicit psychological calculation that the role of prospective feedback and accountability etc. become important in predicting the use of regret minimization as a strategy. Since regret is arrived at via comparison with alternative outcomes, no possibility of feedback would imply no possibility of experiencing regret, which could shift respondents' behaviors towards other strategies.

This commitment to explicit psychological calculation differentiates regret from utility, for which no such commitments are necessary. It is common to observer proposals suggesting direct reward encoding in human observers' brains (Padoa-Schioppa \& Assad, 2006). At a minimum, the idea that utilities may be constructed is not yet consensual in the corresponding literature at the interface between psychology and economics (Slovic, 1995; Srivastava \& Schrater, 2015).

The centrality of explicit calculation for regret is the focus of the work we report in this paper, wherein we sought to characterize the effect of cognitive ease of calculation of regret on decision-makers' meta-decision to use it as a choice strategy. Our hypothesis was that observers would switch away from use of a regret minimization strategy as the cognitive costs of calculating regret increased. To test this hypothesis, we designed a simple choice task wherein expected regret minimization and expected utility maximization yield clearly divergent choice behaviors, and manipulated the choice stimuli to make explicit comparison of items in regret space easier or harder.

We obtained empirical results substantially supporting our hypothesis. Specifically, we found that participants preferred regret minimizing choices when the choice set was a set of monetary labels, but preferred utility maximizing choices when it was a set of product photos, albeit
associated with money labels. A chronometric assessment of difficulty in judging valuation differences between stimuli of the same category was used to establish that regret calculations for the former stimuli category were relatively easier than for the latter. Challenging simple sequential accounts of strategy selection in decision-making, a final experiment demonstrated that decision-makers' stimuluscategory specific bias could be countermanded by changes in the distribution of stimulus valence at the time of presentation. We conclude with a discussion contextualizing our findings within existing formal accounts of strategy selection in decision-making.

## Discriminating between choice strategies

Some econometric research in the past has sought to discriminate between the use of utility maximization and regret minimization strategies by fitting different varieties of discrete choice models to data (Thiene, Boeri \& Chorus, 2012). However, such models have several free parameters and idiosyncrasies in estimation procedures, and their result interpretations are frequently susceptible to validity challenges. To avoid such complications, we sought to design a simple experimental task in which utility maximization and regret minimization would predict clearly divergent choices.


Figure 1: Expected utility (left) and expected regret (right) for nominal x values plotted on the x -axis. A logarithmic form is assumed for the utility function.

This took the form of a choice problem where respondents are told that the correct choice is one of N positive integer-valued options, that each of the options has an equal chance of winning, and that if they guess the correct option, they get the amount of money, or a product of equivalent cost, indicated by the integer value indexing that option ${ }^{1}$.

Our interest was to contrast the relative performance of utility maximization and regret minimization strategies in this setup. Formally, given a set of alternatives X, and some estimate or direct measurement of the utility of alternatives a utility maximizer would select according to the choice rule

[^0]\[

$$
\begin{equation*}
\arg \max _{x} U(x) \tag{1}
\end{equation*}
$$

\]

It is trivial to see that the expected utility maximizing choice in this problem is to always select the option with the highest integer value.

A regret minimizer, on the other hand would calculate the potential regret for choosing each one of the outcomes

$$
\begin{equation*}
R(x)=\left|U^{*}-U(x)\right| \tag{2}
\end{equation*}
$$

where $\mathrm{U}^{*}$ is some counterfactual comparative benchmark utility, and then use the choice rule

$$
\begin{equation*}
\arg \max _{x} R(x) \tag{3}
\end{equation*}
$$

The choice of benchmark utility differentiates regret calculations into different categories. Minimax regret computations take the benchmark utility to be the utility from the best possible outcome (Savage, 1951), and is commonly used in game-theoretic settings to model behavior. Such a criterion is reasonable for when the decision-maker is expected to know the correct option, a common premise in game-theoretic settings. For decisionmakers operating with little domain knowledge, average or expected utility is frequently selected as the benchmark utility, as is common in reinforcement learning settings (Kaelbling, 1996). Since our task falls in the latter category, we use expected utility to perform our regret calculations.

Assuming a linear relationship between utility and regret as defined in Equation (3), we see that the regret minimizing choice in this problem is to pick the option in the middle of the range of available options, calculating $\mathrm{U}(x)$ as the prospective utility of $x$ should it win and treating $U($.$) as a$ logarithmic map of $x$, a classic micro-economic assumption. This pattern is, in fact, inevitable since the benchmark expected utility occurs in the middle of the value range given equi-probable outcomes and draws the regret minimum towards itself. Figure 1 illustrates this intuition quantitatively, showing that prospective regret is lowest when selecting in the middle of the range.

Thus, this simple decision problem potentially gives us a straightforward way of empirically differentiating the use of utility maximizing versus regret minimizing strategies. Assuming even spacing of choice set options, respondents selecting options towards the extreme large values of the offered range are expected utility maximizers, while respondents selecting options in the middle of the offered range are expected regret minimizers.

Given this premise, we next designed a simple experiment to test it. We designed two sets of choice stimuli, one for which regret calculation should be easy, and one for which it should be hard, and asked two different set of respondents to choose between them using the paradigm described above.

## Experiment: easy money and hard pens

Our basic prediction is that decision-makers prefer a regret minimization strategy for option sets wherein comparing the value of options is relatively easy, and prefer utility maximization (or other strategies) when such comparisons are hard. To test this, we designed a between participants' experiment, with one cohort making decisions using a stimuli set that permits easy regret calculations and the other using a stimuli set that does not. As a precursor to this, we ran another study to quantitatively identify which stimuli categories are, respectively, easy and hard for respondents to calculate regret.

## Precursor study

Regardless of whether comparisons are utility functionwise, feature-wise or heuristic-based, it appears natural that the presence of more features should make regret calculations harder. Therefore, we designed choice option sets to have either just one feature (a money amount) or multiple features (money amounts plus other features), corresponding to easy and hard regret calculation settings.

Design. Specifically, we selected two categories of stimuli to test for relative difficulty vis-à-vis a baseline of simple numerical comparisons. These were
(A) two-digit money amounts, and
(B) images of pens, presented alongside their actual market price.
Each participant completed two blocks of 35 trials each for either category of stimulus, with the block presentation order (ABAB/BABA) counter-balanced across participants. Within a block all participants saw a stream of 36 stimuli from a single category ( $\mathrm{ITI}=500 \mathrm{~ms}$ ), and had to successively respond to the cue, "Is this one much better or worse than the last one?" prompting 1 -back comparisons with the stimulus currently on the screen. The sequence of stimuli presentation was pseudo-randomly generated in each category using sampling with replacement from a set of 7 unique stimuli (described in the main experiment for both categories), with the constraint that the new stimulus had to be different from the previous two stimuli in the sequence. "Yes" and "no" responses were coded to the "left" and "right" arrows of a regular QWERTY keyboard. Responses were disabled for the first stimulus in each block since it had no valid comparison. Participants were asked to take as much time as needed to respond, and the trial number within the block was shown alongside the total number of trials in the block on the screen.

Before these four stimuli-specific blocks were presented, participants' response time baselines for numeric distance calculations were established by presenting them with a stream of 36 three digit numbers (ITI $=500 \mathrm{~ms}$ ) sampled from a uniform distribution on $[10,99]$, successively asking the question, "Is this number much larger or smaller than the last one?" The presentation and response interface used for this block was identical to the one used for the subsequent stimuli-based blocks.

Sample and analysis. For this precursor study, we recruited 10 volunteers ( 2 F , age $=24+/-2.3$ years, 0 left-handed) using convenience sampling.

The regret calculation conditions (easy vs. hard), in the form of different stimuli sets, were empirically validated on the premise that the critical step in regret calculation is the utilitarian comparison of the outcome received with an alternative. Adopting a mental chronometric approach, the relative time taken in performing this calculation for different categories of stimuli was used to operationalize our sense of relative difficulty of regret calculations. For all our calculations we report below, we excluded outlier RTs (> 2S.D. from category mean). These constituted $1.5 \%$ of all trials (21 out of 1400 total trials), but occurred primarily in the pens category trials. The exclusion of these outliers in fact deflates the size of primary result we report below. Therefore, we do not report results including them.


Figure 2: Mean Response times for pair-wise difference judgments within different categories of stimuli for all 10 participants of the precursor study. Errors bars represent $+\backslash$ 1 S.D.

Results. Figure 2 displays average response times categorywise, combining trials across participants and category blocks. Clearly, respondents found the monetary comparisons of value approximately as easy as numeric comparisons of magnitude (Cohen's $d=0.40$ ), demonstrating the intuitive mapping of number to value in the monetary domain. Equally clearly, respondents took longer to respond to comparisons involving images of pens alongside their prices (Cohen's $d=1.39$ ), implicating multidimensional considerations in estimating the value of these objects.

Thus the precursor study objectively established that respondents take longer to assess whether two pens offered at different price points are significantly different from each other than to assess this for just two money amounts. Granted the chronometric assumption that RT predicts task difficulty, this result validates our consideration of choice stimuli drawn from the former category as harder than the latter. This distinction, in turn, permits the design and conduct of our main experiment.

## Main study

Design. Volunteers for the main experiment were recruited from the general university population. However, participants from our precursor study were excluded. All consenting volunteers were randomly assigned to easy ( $\mathrm{N}=$ 54 , age $20.4+/-1.4$ years, 31 F ) and hard ( $\mathrm{N}=53$, age 19.8 +/- 1.0 years, 25 F ) regret calculation conditions respectively.

Both sets of respondents participated in the experiment in a classroom setting separated spatially from each other, transmitting their responses via text messages. The easy group respondents were presented with the following instructions, "Consider this hypothetical scenario. I have a bowl of seven paper tokens, each one with one of the first seven multiples of five written on them. Every number is written on at least one token, and no token has more than one number on it. At the end of the class, I will draw a token and whoever can text me (response number) the number on the token I will draw will win the amount of money written on that token."


Figure 3: Choice stimuli presented to respondents in the hard condition. Numbers in parentheses represent pen codes. Money amounts are true prices of the corresponding pens. Pens are arranged in randomized order with respect to money amounts to prevent positional bias in responses

The hard group respondents were presented with the visual display shown in Figure 3 accompanied by the instructions, "Consider this hypothetical scenario. I have a bowl of seven paper tokens, each one with a number between 1 and 7 written on it. Every number is written on at least one token, and no token has more than one number on it. At the end of the class, I will draw a token and whoever can text me (response number) the number on the token I will draw will win a pen of the type listed under that number on this display."

Analysis and results. Figure 3 summarizes the responses from both groups of respondents as a histogram of the number of respondents that selected each response option. The difference between the response patterns is visually apparent in the modes of the two distributions in Figure 3,
and a two-sample T-test of the individual responses from the two cohorts also indicates a significant difference $\left(\mathrm{t}_{105}=\right.$ 2.18, p = 0.03). An effect size calculation yielded a Cohen's $d$ of 0.41 , again consistent with a significant difference between the two response patterns.

A comparison with the predictions from Figure 1 clearly suggests that respondents from the easy group, who were significantly biased towards responding in the middle of the proffered range, were likely using a regret minimization strategy, whereas respondents from the hard group, who preferred the pricier pens, were likely using a utility maximization strategy.


Figure 4: Histogram of respondents' selections for choices where regret calculation is designed to be (left) easy and (right) hard.

This finding is not easily explicable by alternative hypotheses. Previous theoretical proposals have suggested that respondents prefer to decide based on prospective regret when choices are difficult or consequential (Zeelenberg \& Pieters, 2007). If anything, it appears intuitive that choosing in pen space is more difficult than choosing in money space. Results from our precursor study establish, at the very least, that estimating value differences between pens in our display is harder than estimating value differences between money amounts. If the pens are harder to choose from, then Zeelenberg \& Pieters (2007) would predict the opposite pattern of results than what is seen. Similarly, arguments explaining regret minimization being preferred in unfamiliar domains should also predict it being used when selecting between pens than between money amounts, since choosing between money amounts is unlikely to be more unfamiliar than choosing between idiosyncratic stimuli like pens. Thus, this result appears to clearly favor an ease of calculation explanation for preferring a prospective regret minimization strategy.

## Input or enabler?

While the difference in responding elicited by our manipulation does suggest a role for the ease of regret calculation entering into respondents' decision about which strategy to use, it does not clarify how this variable enters this reasoning.
We conducted a variant of the original experiment to differentiate between two potential roles for this cognitive effort variable: (i) as an input to hierarchical decision
process, where the strategy is selected first, then implemented, followed further by assimilation of feedback, or (ii) as a mechanistic enabler, in the sense that quicker regret calculation makes results from the use of a regret minimization strategy available sooner to participants, and hence more likely to be used.
As we discuss further below, the first possibility would fit this cognitive effort variable within formal hierarchical models of strategy selection and learning, such as Rieskamp \& Otto's influential SSL model (Rieskamp \& Otto, 2006). The latter would be more compatible with heuristic accounts of the effect of availability and accessibility on decisionmaking (Carroll, 1978), which are yet to be successfully formalized to the same extent (Gigerenzer \& Gaissmaier, 2011).

Design. If strategy selection precedes outcome evaluation, then we would expect changes in the range of outcomes used for our decision problem to not affect the choice of decision strategy. Conversely, if changing the range of outcomes for the decision problem reveals differences in the pattern of responding, it is clear that some aspects of outcome evaluation must precede the decision of which strategy to use.

To test this hypothesis, we again used a between-subjects design. The decision problem and setup was identical to the one used in the easy condition of the main experiment, with two different groups of respondents making choices using two different sets of money amounts. The first set used the same stimuli as the original experiment. The second set used the stimulus set $\{5,10,15,20,25,30,105\}$, replacing the largest stimulus in the original set with a much larger value. All participants received the same instructions as in the main experiment's easy condition in a classroom setting, and transmitted their selections using text messaging as before.

Sample. Volunteers for the experiment were recruited from the general university population. We screened the recruited sample for previous participation in either our precursor study or the main experiment. A total of 90 participants (23F, Age $=20.3+\backslash-1.8$ years) were finally selected for participation in the experiment, and randomly assigned to two equally-sized groups for this study.


Figure 5: Histogram of responses for groups responding to (left) original stimuli set and (right) changed stimuli set

Analysis and results. As is evident from Figure 5, the response patterns in both groups were starkly different. A two sample T -test for the individual responses returned strongly statistically significant $\mathrm{t}_{88}=6.00, \mathrm{p}<10-6$ and an effect size calculation yielded a large effect (Cohen's d = 1.08). Further, the response pattern elicited from the 45 participants who were presented with the same stimuli as in the easy condition in the main experiment were not statistically different from the 54 participants' responses obtained during the former experiment (two sample t-test p $=0.40)$, suggesting that the original result was robust.
This result shows that changing the set of choice options by adding an extremely valuable alternative makes respondents substantially more likely to prefer the expected utility maximizing strategy, suggesting that a sequential view of strategy selection followed by application cannot faithfully reflect how participants use information available from the choice set before making their decision. Thus, the present evidence suggests that the expected different costs of regret computation for different stimuli sets does not enter explicitly into participants' strategy-selection calculations, but rather enables regret-based determinations to be emitted preferentially by virtue of being generated quicker, in line with the bag of heuristics view of decisionmaking strategies (Gigerenzer \& Gaissmaier, 2011). However, we discuss below how our findings could potentially be reconciled with a hierarchical view of strategy selection further below.

## Discussion

Summary of results. In this paper, we have proposed a novel characterization of when human decision-makers are likely to prefer minimizing prospective regret over alternative decision-making strategies like expected utility maximization. Our proposal is that decision-makers prefer to minimize regret when the cognitive cost of calculating regret is low, and switch to alternative decision strategies when this cost is high.
To test this hypothesis, we designed a simple decision problem which permits a clear empirical differentiation between the use of either of these two decision-making strategies. We conducted a chronometric assessment of two stimuli sets, for which relative value judgments had distinctly different difficulty levels. Although other measures of effort have been proposed in the literature, drawing upon information-theoretic considerations (Huber, 1980), the validity of these measures is ultimately assessed using reaction time data (Johnson \& Payne, 1985). Thus, while alternative operationalizations of effort are certainly possible, our response time-based definition appears reasonable.
Using this observation to establish the relative difficulty of regret calculations using options selected from these two stimuli sets, we asked two separate groups of participants to make decisions that were formally identical, except for the stimuli identity difference. We found that the pattern of responses for choices made using stimuli that were hard to
evaluate comparatively was more consistent with the use of an expected utility maximization strategy, whereas for stimuli that were easier to compare, the pattern of responses was more consistent with the use of an expected regret minimization strategy. It is, of course, impossible to verify that these were the only two strategies possible for participants to use in the choice problem. Ad hoc heuristic approaches such as 'bias towards the middle of the range' etc could, in principle, be potentially confounded with the regret minimizing predictions for this choice problem. Such ad hoc proposals, however, are not parsimonious, in the sense that they fail to explain the shift to expected utility maximization for the same choice problem using different stimuli, whereas the cost of calculation explanation does.

A more significant question is how well the result demonstrated in this somewhat arbitrary choice problem generalize to richer experimental settings and real-world decisions. We consider this an important consideration for future work.

Related work. There is a large literature on strategy selection, anchored in contemporary times by Rieskamp \& Otto's powerful SSL theory (Rieskamp \& Otto, 2006). The basic outline of this theory is that observers select a strategy to tackle each instance of a decision problem stochastically, guided by their preference for each of the possible strategies. This strategy-preference in SSL has three components, (i) the maximum reward possible in a trial, (ii) an initial strategy-specific preference, and (iii) a learningbased association of strategy to the choice problem, based on the long-run trend of the use of that strategy resulting in a higher reward. The basic intuition underpinning SSL is that observers adapt to choice contexts by gradually learning to prefer strategies that prove more rewarding in them. Notably, Rieskamp \& Otto (2006) explicitly consider the possibility that the cognitive costs of applying a strategy may enter observers' calculations for strategy preference. However, how such strategy-specific costs would enter their model's calculation has remained an open question.
The results in this paper provide useful constraints on the potential development of such a cost-sensitive model. Our main experiment strongly suggests a role for cognitive cost of applying a strategy in determining observers' preference for it. A naïve approach might be to subtract some notional cost of calculation from the reward term in the SSL prior on strategy preference. However, our follow-up experiment demonstrates that strategy preference can be affected by complex informational aspects of the choice problem, such as the distribution of options in value space.

Such a complex interaction does not appear to be possible in the baseline two-step algorithmic specification of SSL, wherein first the strategy is selected based on existing strategy preferences, and then information from the current trial updates the strategy preferences. We conjecture that a drift-diffusion based (Ratcliff \& McKoon, 2008) extension of the SSL model, wherein the evidence for the utility of options accumulates competitively and becomes available to
assist in strategy evaluation depending on how soon this competition terminates, could accommodate our results.

## References

Boeri, M., Scarpa, R., \& Chorus, C. G. (2014). Stated choices and benefit estimates in the context of traffic calming schemes: Utility maximization, regret minimization, or both?. Transportation research part A: policy and practice, 61, 121135.

Carroll, J. S. (1978). The effect of imagining an event on expectations for the event: An interpretation in terms of the availability heuristic. Journal of experimental social psychology, 14(1), 88-96.
Connolly, T., \& Reb, J. (2005). Regret in cancer-related decisions. Health Psychology, 24(4S), S29.
Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., \& Dolan, R. J. (2005). Regret and its avoidance: a neuroimaging study of choice behavior. Nature neuroscience, 8(9), 1255.
Gigerenzer, G., \& Gaissmaier, W. (2011). Heuristic decision making. Annual review of psychology, 62, 451-482.
Huber, O. (1980). The influence of some task variables on cognitive operations in an information-processing decision model. Acta Psychologica, 45(1-3), 187-196.
Johnson, E. J., \& Payne, J. W. (1985). Effort and accuracy in choice. Management science, 31(4), 395-414.
Kaelbling, L. P., Littman, M. L., \& Moore, A. W. (1996). Reinforcement learning: A survey. Journal of artificial intelligence research, 4, 237-285.
Loomes, G., \& Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. The economic journal, 92(368), 805-824.
Padoa-Schioppa, C., \& Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. Nature, 441(7090), 223.

Ratcliff, R., \& McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural computation, 20(4), 873-922.
Rieskamp, J., \& Otto, P. E. (2006). SSL: a theory of how people learn to select strategies. Journal of Experimental Psychology: General, 135(2), 207.
Sarver, T. (2008). Anticipating regret: Why fewer options may be better. Econometrica, 76(2), 263-305.
Savage, L. J. (1951). The theory of statistical decision. Journal of the American Statistical association, 46(253), 55-67.
Slovic, P. (1995). The construction of preference. American psychologist, 50(5), 364.
Small, K. A., \& Rosen, H. S. (1981). Applied welfare economics with discrete choice models. Econometrica: Journal of the Econometric Society, 105-130.
Srivastava, N., \& Schrater, P. (2015). Learning what to want: context-sensitive preference learning. PloS One, 10(10), e0141129.
Thiene, M., Boeri, M., \& Chorus, C. G. (2012). Random regret minimization: exploration of a new choice model for environmental and resource economics. Environmental and resource economics, 51(3), 413-429.
Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making. Journal of behavioral decision making, 12(2), 93-106.
Zeelenberg, M., \& Pieters, R. (2007). A theory of regret regulation 1.0. Journal of Consumer psychology, 17(1), 3-18.


[^0]:    1 The inspiration for this problem design is drawn from an unpublished draft by Oleg Urminsky \& Adele Yang, which in turn derived this problem from a common radio station contest - the jackpot guessing game.

