Groups are Better than Individuals at Solving Optimum Stopping Problems

Pritam Laskar (pritaml20@iitk.ac.in)
Department of Cognitive Science
IIT Kanpur, India

Nisheeth Srivastava (nsrivast@iitk.ac.in)
Department of Computer Science and Engineering
IIT Kanpur, India

Abstract

Many real-life decisions, such as booking a vacation or selecting a partner, involve relatively cost-free sampling of options up to a terminal decision point, beyond which the choice becomes costly to reverse. Such problems can be formulated as optimal stopping problems (OSPs), such as the famous secretary problem. Although human behavior on optimal stopping problems has been studied extensively, much of the literature has focused on the behavior of individual decision-makers operating using a binary payoff function. In this study, we use an OSP with a continuous payoff function to study how individuals' decisions differ from the collective decision of groups of three members working together. An independent threshold model offered the best explanation for the behavior of both individuals and groups. We found groups performed significantly better than individuals, with individuals consistently waiting too long to make a choice relative to the optimal strategy. Groups are also more decisive in following their internal thresholds, which are also different than a simple average of member thresholds. Finally, we also found a lack of long-term learning in OSPs for groups, a trend previously documented in individuals.

Keywords: Decision Making, Optimal Stopping Problem, Group Decision Making, Bayesian Modelling, Threshold Models

Introduction

Consider this scenario - you want to buy a gift online for a friend's birthday and you must do it in 10 days to get it delivered on time. You have a particular gift in mind and your goal is to buy the item at the cheapest possible price. Assuming price fluctuations are a somewhat random process, this may be formulated as minimizing the expected price paid taking such fluctuations into account. You check the item's price every day and if it is too high you wait for the next day. Once you buy the item the search stops and the decision is irreversible (let's assume there are no refunds/returns). However, since the gift must be bought within 10 days, if you still skip buying on day 9 you are forced to buy at whatever the selling price is on day 10. Note that the reward is not all or nothing - even if you end up buying the gift at the second-cheapest price, you may be satisfied with the outcome. This scenario is an example of a general class of optimal stopping problems (OSPs), which have been mathematically studied extensively beginning with Gilbert and Mosteller's (1966/2006) work.

It turns out that if certain assumptions are made - for instance that the prices follow a given distribution $f(x)$ and are independently sampled on each day - a mathematically optimal stopping rule can be found. In the above scenario, the optimal scenario is threshold based. What this means is that for days 1 through 10 there are increasing threshold prices $p_1, p_2, \dots p_9, p_{10}$ and the optimal strategy is to buy on the first day i when the asking price is less than p_i (Gilbert and Mosteller, 2006). These thresholds are calculated recursively - p_{10} is $+\infty$, as you are forced to buy no matter what price and the expected price paid in this case equals the mean of the distribution, μ. p_9 equals the expected price paid on skipping on day 9 and so is μ. p_8 is set equal to the expected price paid when one skips on day 8, and so on. In general:

$$p_i = \int_{-\infty}^{p_{i+1}} x f(x) \mathrm{d}x \; + \; p_{i+1} \cdot \int_{p_{i+1}}^{+\infty} f(x) \mathrm{d}x, \quad \forall i = 8, 7, \dots 1$$

Kinds of Stopping Problems

Rank vs. Full-information: OSPs can broadly be categorized into rank information and full information problems. In rank information problems, like the famous secretary problem, only the relative ranks of each option with respect to all seen options are known. Interviewing the 5$^{\text{th}}$ candidate among 20 total candidates, you only know that he is second best so far. In full information problems, the cardinal value of each option is explicitly given. This is the case in the scenario presented above where the price of the gift tells you everything about the option of purchasing on a given day. Both types of stopping problems can be useful models for different real-world scenarios. Looking for a mate or a good nesting spot for a bird may be an irreversible decision with no clear numerical value tagged to each option. On the other hand, many problems in the modern world such as when to buy a stock are full-information scenarios. The form of the solution for an OSP crucially depends on the information provided. Optimal solutions to secretary-type problems (with binary payoff) involve rejecting the first $n\%$ candidates ($n \to 1/e \approx 37\%$, as total candidates $N \to \infty$) whereas optimal solutions to full-information problems are usually threshold-based regardless of the payoff function (Ferguson, 1989).

Payoff Function: The other important factor in determining the optimal solution is how rewards are given once a choice is made. Rewards may be binary, i.e., only given if one selects the best alternative, or continuous. For the gift-buying scenario above, this difference translates into trying to maximize the probability of buying the gift at the lowest price vs. minimizing the expected price paid for the gift, which is how

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society.*

we framed the problem. Real-life scenarios such as choosing a job or buying shares in a company rarely give binary returns - *some* payoff is obtained even if we don't end up with the best job, or time our entry into the market perfectly. Therefore continuous payoff functions are potentially a more realistic model for studying OSPs vis-a-vis humans. Nonetheless and somewhat surprisingly, earlier behavioural research on OSPs has dealt almost exclusively with binary payoffs (Baumann et al., 2020).

Previous Work

A number of previous studies on OSPs have dealt with individual behaviour - both on rank-information secretary problems (Seale and Rapoport, 1997; Seale and Rapoport, 2000; Campbell and Lee, 2006) and full information problems (Baumann et al., 2018; Guan and Lee, 2018; Lee, 2006). Compared to the mathematically optimal solutions, individual performances have been repeatedly found to be worse, in ways that we characterize more precisely below. This may be due to individuals following a different heuristic or form of decision rule. Or perhaps, individuals do use the same form as the optimal solution but use sub-optimal cutoffs/thresholds. Moreover the difference is not always straightforward to interpret. For instance, Seale and Rapoport (1997,2000) showed through simulations that in the original secretary problem an alternate heuristic - a successive non-candidate rule, where the first top ranked interviewee seen after *k* successive non-top ranked interviewees is selected - can be up to 98% efficient as the optimal cut-off rule. However, some robust results have been found.

- **Lack of long-term learning:** On both rank and full information problems individual performance shows no long-term learning effects across blocks. This has been observed in rank information OSPs (Seale and Rapoport, 2000) even when using monetary feedback (Campbell and Lee, 2006). A similar effect is seen in full-information problems as well (Baumann et al., 2020; Lee, 2006).
- **Threshold models describe full-information OSP behaviour:** Individuals seem to use threshold-based strategies for their choices in full-information problems. This holds for both binary payoffs (Guan et al., 2014, Guan et al., 2015, Lee, 2006, Lee and Courey, 2021) and continuous (Baumann et al., 2020) payoffs. However a number of different threshold models have been proposed - some simpler than others. Recently, Baumann et al. (2020) compared 3 different threshold models for a continuous payoff function and found a linear threshold model to best describe individual behaviour with continuous payoffs. Recently, Lee and Courey (2021) found a fixed-then-linear threshold described behaviour in a changing environment (where the distribution of values changes as the sequence progresses) well.
- **Direction of sub-optimality:** In OSPs, sub-optimal behavior may arise due to stopping too early as well as not stopping soon enough. A robust trend seen across a variety of rank and full-information OSPs is that individuals tend to stop too early in case of binary payoffs (Baumann et al., 2018, Guan et al., 2014, Guan and Lee, 2018, Seale and Rapoport, 1997). In case of continuous payoff, Baumann et al. (2020, SI Appendix) using the same payoff function and price-minimization task as ours and found individual price thresholds to be higher (less strict) than optimum initially with most options left and lower (stricter) than optimum later on when few options remain.

Group Performance in Stopping Problems Real-world decisions in organizations, governments and even families are often made not by individuals but groups of individuals. Many studies in the past have found consistent differences between individual and group decision making. Group decisions in games have been found to be much closer to game theoretic rationality than individual ones (Kugler et al., 2012). Teams, by which we mean groups with no conflict of interest and a common goal, also behave differently from individuals when it comes to uncertain and temporal decisions (Kocher et al., 2020). Given these observations, it is possible that groups behave differently in OSPs than individuals.

An important work for group behavior in OSPs comes from a study by Lee and Paradowski (2006), who had individuals and groups work out a full-information OSP with a binary payoff. In their study the goal was to select the maximum of 5 integers drawn randomly between 0-100 without replacement. Groups were made of 5 members who did not communicate face to face but through a computer interface. On being presented a number members made an initial choice of whether or not to chose the current number. This was seen by all other group members. Members then had a chance to revise their decision, before the final group decision was determined through a rule. They considered three decision rules - majority, consensus and leadership. Using a signal detection theory framework, possible due to binary payoff, they classified response on each presented number as correct or incorrect. They found higher d' values for consensus and leadership conditions resulting in better performance as compared to individuals. These two conditions also had higher bias c and so groups did not stop as early as individuals. Recent work by Thomas et al. (2021) provides potential insight into how groups may be better in OSPs. By aggregating the inferred individual thresholds, they formed a model crowd whose simulated behaviour matched the behavioural crowd and whose accuracy was higher than the participant average.

Research Questions

In the following, we will use the terms 'teams' and 'groups' interchangeably as there is no conflict of interest among group members in our experiment. Two primary questions motivate the current study, in light of the gaps in the literature documented above:

1. *Do small teams working on a full-information OSP with a continuous payoff function perform better/differently as*

compared to individuals?

We predicted groups would perform better given previous results by Lee and Paradowski (2007).

2. *What sort of threshold model best describes individual and team behaviour in a continuous payoff OSP?*

   We used Bayesian model comparisons to test three threshold models - the Linear Threshold Model (LTM), Independent Threshold Model (ITM) and a Biased Optimum Model (BOM) with individual and team behavioural data. These models are described in a later section.

To answer these questions we used the same scenario and payoff function as in Baumann et al. (2020), with slight modifications as described below.

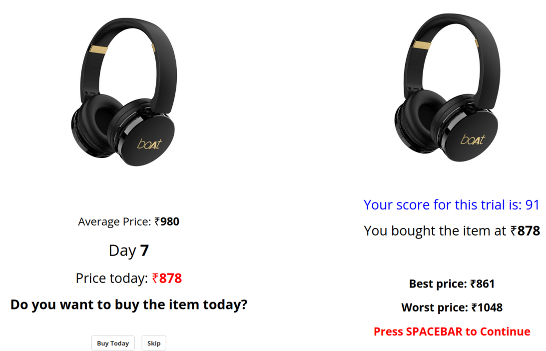

Figure 1: Screenshot of a Scored Trial

Experiment

An IRB approved the protocol for our study. Our experiment is hosted online but was conducted in-person under supervision. 27 individuals grouped into teams of 3 members each participated in the study. Each participant performed the same task twice - once alone in the *individual phase* and once as part of the 3 member team in the *group phase* (counterbalanced). Participants were university students between ages 18-30. They were all compensated ₹100 for their participation in the study, which took around 75 minutes for both the phases. Importantly, teams were not created randomly and participants often brought in people they knew well as team members. As a result, in 6 of the 9 groups the participants themselves were friends or acquaintances.

Task

The goal of the experimental task was to buy an item at the least possible price within 10 simulated days. There were 5 real-world products in total (including headphones, t-shirts, microwave ovens, suitcases and shoes) available for purchase off e-commerce platforms. Prices for an item were drawn independently on each day from a normal distribution whose mean μ was equal to the item's online selling price. Each of the 5 products was repeated twice - a high-variance condition with double the standard deviation σ as the low variance condition for a total of 10 item blocks. The ratio $\frac{\sigma}{\mu}$ was approximately 0.05 in the low variance condition and around 0.1 for the high-variance condition for each item. Each item block had 10 trials - 2 practice trials followed by 8 scored trials, giving a total of 80 scored trials in total for each individual and group. Items were chosen to reflect a large range of price values (from ₹300 to ₹6000).

The individual or group decision maker (DM) always knew the mean of the distribution which was displayed on-screen. However they were not told the variance or the type of distribution (Gaussian) the prices were drawn from. They learned about the distribution *through acquaintance* during the trials. However, participants were specifically informed of the independence in prices drawn - something not seen in the real world where daily online prices are correlated and display predictable trends approaching holiday sales. On days 1 through 9 DMs had the option to either buy the item at the displayed price or skip to the next day. If an item was not bought by day 10, DMs were forced to buy at the quoted price. Scored trials ended as soon as an item was bought. However each item block began with 2 practice trials where DMs had to manually go through prices on all the 10 days. This was so that every participant and group began the scored trials after seeing exactly 20 sample prices.

For each practice and scored trial an integer score between 0-100 was assigned to the trial performance based on the following formula (Baumann et al., 2020):

$$\text{Score} = 100 \times \left(\frac{P_{max} - P_{choice}}{P_{max} - P_{min}} \right) \tag{1}$$

Here P_{max}, P_{min} denote the highest and the lowest of the 10 prices in a trial, respectively. P_{choice} denotes the price the DM bought (or was forced to buy) at. Thus, if on a trial the decision maker (DM) managed to buy the item at the cheapest price she received 100 points and 0 points if she bought it at the highest price. Equation 1 defines a continuous payoff function for our optimal stopping task. Participants were shown their trial score along with these three prices after each trial (see Figure 1). To motivate DMs to perform their best, participants were told a part of their compensation would depend on their average score across all the scored trials - both in individual and group phases, equally weighed.

Group Phase Inputs in the task were made using a mouse and a keyboard. This is straightforward in the individual phase where participants did the experiment alone and were not allowed to interact with anyone. In the group phase the 3 team members were seated together in front of a computer and followed a protocol. On each of the 10 item blocks, one member was in-charge of making the inputs. Although advised to use a majority decision rule, groups were free to use any process to make the final choice. The only requirement was that members discuss among themselves before the person in-charge makes the final input. An experimental supervisor was always present throughout to ensure the protocol was being followed. All the groups ended up either following a consensus approach (where, all members were okay with going ahead with a response, even if 1-2 members disagreed) or a majority rule (where the choice getting 2 or more votes was chosen). Importantly, the person in-charge was rotated for each item block so that one member was in-charge of 4 items and the other two members for 3 items each.

Modeling Optimal Stopping

As we describe above, the mathematically optimal solution to our continuous payoff OSP is threshold based. These thresholds increase with the day count. This makes intuitive sense - when you have 10 days left to buy the item you can afford to have stricter standards. As the days run out you are willing to accept paying higher prices for an item to avoid the risk of being left with nothing. Following Baumann et al. (2020) we model the decision maker's (whether individual or group) choice of either buying or skipping at asking price x_i using an internal threshold T_i on day i as a Bernoulli random variable with probability of buying p_i given by:

$$p_i = \frac{1}{1+e^{\beta(x_i - T_i)}} \tag{2}$$

Here, $\beta \geq 0$ is a parameter controlling how decisively the DM follows their internal threshold. When $x_i < T_i$, the DM is more likely to buy the item and vice versa. As $\beta \to \infty$, the agent begins to always buy when price is lower than threshold and never buy when $x_i > T_i$. $\beta = 0$ reflects buying with probability 0.5 regardless of price.

The models we consider differ in the other parameter in Equation 2 - the thresholds T_1 through T_9 [1]. Since we have used ten items whose prices come from ten different normal distributions, we used normalized prices and aggregated responses from all ten item blocks to calculate posterior distributions for all model parameters.

1. **Linear Threshold Model (LTM):** T_i's increase linearly with day i with common difference δ, giving:

   $$T_i = T_1 + \delta \cdot i \tag{3}$$

   This model has 3 parameters - β, δ and T_1 and does not include the optimal strategy for prices coming from a normal distribution. We used uniform priors $\beta \sim U(0.1, 10)$, $\delta \sim U(0, 1/3)$ and $T_1 \sim U(-3, 0)$ for both individual and group data.

2. **Independent Threshold Model (ITM):** No assumptions are made about T_i's. This is model encompasses the optimal solution, any LTM model and has 10 parameters - β and 9 thresholds $T_1, \ldots T_9$. Priors used were $\beta \sim U(0.1, 10)$, $T_i \sim U(-3, 0.5)$ for $i = 1, \ldots 9$, and were the same for individual and group data.

3. **Bias from Optimum Model (BOM):** T_i differs from optimal threshold p_i through:

   $$T_i = p_i + \gamma + \alpha \cdot i \tag{4}$$

   This model is taken from Guan et al. (2015) and has 3 parameters - β, γ and α. γ gives a fixed bias from the optimum threshold which may increase or decrease as the sequence progress, depending on α. Again we used the same priors for individual and group data - $\beta \sim U(0.1, 10)$, $\gamma \sim U(-3, 3)$ and $\alpha \sim U(-3, 3)$.

[1] $T_{10} = +\infty$ as the DM is forced to buy on day 10

Baumann et al. (2020) found the LTM to be almost as good as the ITM, while BOM gave a poor fit to individual behavioural data. Since LTM is more parsimonious, they inferred that individuals behave as if using simpler linear thresholds. We must add that recently Lee and Courey (2021) found a fixed-then-linear threshold model accounted well for experimental data in a non-stationary setting. We have excluded this from our current analysis, ours being a stationary OSP focused on differences between individuals and group behaviour.

Results

Learning Across Trials

In order to detect outliers, we assumed the mean score for groups and individuals to be coming from separate Student's-t distributions and rejected values lying 3 standard deviations away from the means. This resulted in rejection of a single individual from a sample size of 27 from further analysis. No outliers were detected among the 9 groups.

Group vs. Individual Performance

On scored trials, average groups scores ($n = 9$, $M = 87.82$, $SD = 1.23$) are greater than average individual scores ($n = 26$, $M = 85.62$, $SD = 3.10$). A Bayesian independent sample T-test provided *some* evidence for the alternate hypothesis that group scores are greater than individual scores ($BF_{+0} =$ 2.923, using default Cauchy prior with scale = 0.707 for effect size δ in JASP).

We also found that the variance in group scores is lower compared to individuals scores, supported by Levene's test for unequal variances ($F(1, 33) = 4.79$, $p = 0.036$). This suggests that groups' decisions are more similar to each other than individuals'. Interestingly and relatedly, Shupp and Williams (2008) compared risk-aversion behaviour of individuals with three-person groups and found variance in CERs (Certainty Equivalent Ratios) to be less for groups, especially in high-risk lotteries (with low chance of winning). They also found groups to be more risk-averse in high-risk lotteries and more risk-seeking in low-risk lotteries as compared to individuals.

Figure 2 shows both individual and group performance across 4 quarters. Since there are 80 scored trials in total, each quarter is the average of 20 scored trials. Visually it seems as though both individual and group performances improve initially from quarter 1 to quarter 2 but then saturate. We used JASP to perform a Bayesian Repeated measures ANOVA using quarter scores and found very strong evidence for an effect of quarter on individual scores ($BF_{10} = 38.676$) and moderate evidence for the same in groups ($BF_{10} = 7.717$).

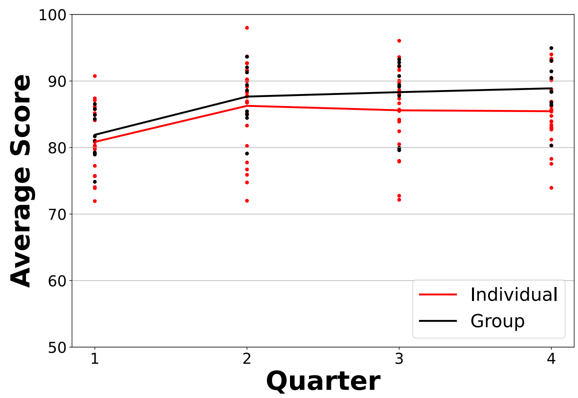

Figure 2: Performance across quarters

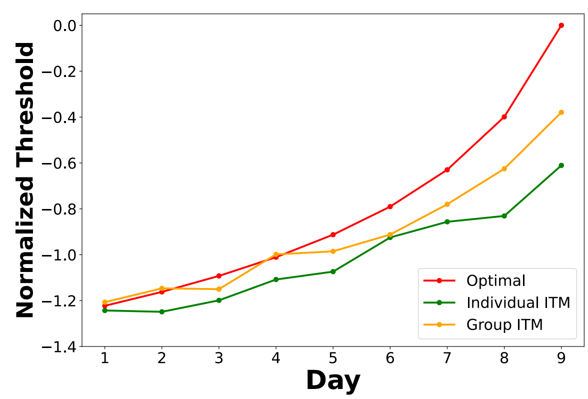

Figure 3: Aggregated individual and group thresholds

Bayesian post-hoc tests confirmed that for individuals only quarter 1 scores were different from quarters 2 through 4 (corrected posterior odds ≥ 5.855 for all 3 comparisons). Comparing other quarters pairwise provided strong evidence for the null that they did not differ (corr. posterior odds ≤ 0.097). For groups too, all pairwise comparisons for quarter 2 through 4 performances provided moderate evidence for the null (corrected posterior odds ≤ 0.154). Post hoc comparisons only found anecdotal evidence for lower quarter 1 scores once corrected for multiple comparisons (corr. posterior odds ≥ 1.252, uncorrected $BF_{10} \geq 3.021$), hinting that a larger sample size for groups is warranted in the future. These results confirm and extend to groups previous results which showed that individual performance quickly saturates and shows no long-term learning towards the optimal value.

Model Comparisons

We used the PyMC3 package to generate posterior distributions using the NUTS sampler (Salvatier et al., 2016). The Arviz package was used to compare models based on their expected log pointwise predictive density (ELPD); estimated using Pareto smoothed importance sampling based leave-one-out cross-validation (LOO) (Vehtari et al., 2017).

Table 1: Model Comparison for Individual Data

Model	ELPD Diff.	SE Diff.	Weight	BF_{ITM}
ITM	-	-	0.971	1
LTM	12.76	6.32	2.87e-2	33.8
BOM	56.07	11.45	6.64e-10	> 1000

Table 2: Model Comparison for Group Data

Model	ELPD Diff.	SE Diff.	Weight	BF_{ITM}
ITM	-	-	0.798	1
BOM	4.60	5.02	0.197	4.05
LTM	19.12	7.28	0.001	> 100

Tables 1 and 2 list the three models in decreasing order of posterior weights for individuals and groups, respectively. It lists the ELPD difference of a model from the best model, which is the ITM in both cases, along with the standard error of this value. The final column lists the Bayes Factor of the ITM with respect to other models. Clearly individual behaviour is best described by presuming no relationships on the internal thresholds used. For groups, the ITM moderately outperforms the BOM model ($BF_{ITM} = 4.05$). In contrast to Baumann et al. (2020), who preferred the LTM in all three experiments, we found the LTM to be very poor in accounting for individual data. However, their interpretation was based on the fact that LTM posterior-predictive p-values was not too worse off from that of ITM. Since LTM has fewer parameters, they preferred it over ITM. We compared the models directly and found the ITM to be superior based on the posterior model weights and associated Bayes factors, where the ITM was automatically penalized for extra parameters.

Direction of Suboptimality

Figure 3 plots the optimal thresholds and the normalized [2] mean group and individual posterior thresholds aggregated over all ten item blocks per day, using the ITM. This makes the source of improved group performance very clear. Individual thresholds are consistently lower than optimal, showing that in our task individual standards were 'too strict' regardless of the number of options remaining. This is in stark contrast to previous findings in full-information problems with binary payoff (Baumann et al., 2018, Guan et al., 2014, Guan and Lee, 2018) and continuous payoff (Baumann et al., 2020) where DMs stopped too early. Groups, in contrast, were more comfortable buying at higher prices and appropriately changed expectations as options dwindled. From the especially low individual thresholds on later days, it seems individuals keep choosing a higher-risk gamble for a good price

[2] Normalized thresholds are simply the Z-scores of the actual thresholds, calculated after combining the normalized prices from all ten item blocks

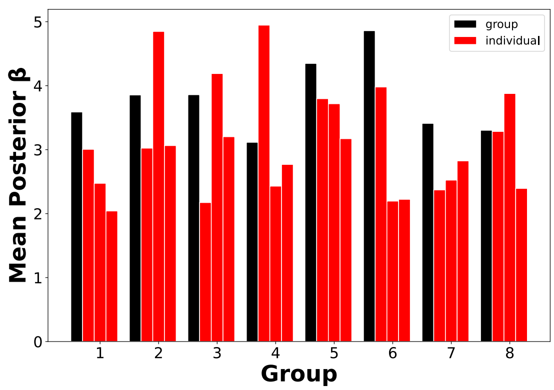

Figure 4: Group and member β

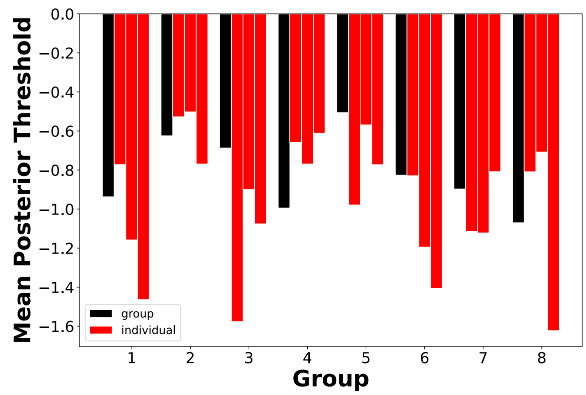

Figure 5: Group and member thresholds for Day 7

and do not appropriately lower their expectations. This pattern of behavior is consistent with the findings of Shupp and Williams (2008) showing greater risk-seeking in individuals as compared to interacting 3-member teams.

Individual to Group Behaviour

Fitting data generated by each individual and group to the ITM, we found decisiveness β and thresholds T_is for each DM. Figures 4 and 5 give bar plots of β and T_7 for 8 groups [3] along with their three members. The plots indicate that groups are much more decisive than their members are individually as well as have higher (less negative) thresholds. Using JASP, we performed Bayesian paired T-tests confirming this - $BF_{+0} = 5.570$ for H_a that groups have greater β than the average of members and $BF_{+0} > 3500$ for H_a that group thresholds (pooled across 9 days) are greater (less-negative) than the average of members. These findings suggest that something more than a simple aggregation of independent member decisions is going on here.

Conclusion

We found significant differences between group and individual behaviour on a full-information OSP with continuous payoffs. Interacting 3-member teams perform significantly better and are less variable than individuals in knowing when to terminate search and make a selection. Our findings confirm previous observations of quick saturation of individual performance at sub-optimal levels in OSPs, extending this result to groups as well.

Interestingly, the nature of individual sub-optimality observed here is inverted from the pattern seen previously for full-information OSPs with binary payoffs Baumann et al., 2020. Our results also differ in some aspects from those documented in an earlier study with continuous payoffs Baumann et al. (2020), who found a change in the nature of suboptimality in individuals across their experiment trial days. Their participants stopping too early initially and then too late, failing to reduce standards fast enough as options diminished. This was the case across all their experiments, including one (Experiment 3 in Baumann et al. (2020)) that used exactly the same protocol as our experiment for individual subjects. In contrast to their findings, our individual participants stopped too late rather than too early throughout the course of the trial instead of only doing so towards the end. Reassuringly, consistent with Baumann et al. (2020), we too found that the difference from optimum for individuals worsened over time, as they failed to decrease standards fast enough. Groups also stopped too late, but this bias was much less pronounced, especially with more options remaining. We also find that both group and individual behaviour are best described by the ITM, presuming no relationship among the internal thresholds, in contrast to Baumann et al. (2020). These differences cannot be reconciled based on existing evidence, and must be investigated in future work.

Our analysis also found that group behaviour was more complex just an aggregation of the members' individual decisions (Thomas et al., 2021). In the ITM model, this is reflected in how group thresholds were different than the average member threshold. Interestingly, groups were more decisive than individuals as well in following their internal thresholds than individual members. Future work can examine how well models of individual threshold adjustment can accommodate these results (Srivastava et al., 2016).

An important limitation of our current study is that it does not engage very much with the fact that groups take much longer to decide on their choices than individuals. This is by design as members were instructed to discuss before making a choice. Future work could control for this by forcing individuals to wait for a comparable time before making their choice, and/or monitor the process of evidence integration within the groups so that the mechanisms through which preferences and differences in individual members amalgamate into the final decision can be better understood.

[3] We a-priori excluded the group having an outlier participant from the analysis in this section.

References

Baumann, C., Singmann, H., Gershman, S. J., & von Helversen, B. (2020). A linear threshold model for optimal stopping behavior. *Proceedings of the National Academy of Sciences*, *117*(23).

Baumann, C., Singmann, H., Kaxiras, V. E., Gershman, S., & von Helversen, B. (2018). Explaining human decision making in optimal stopping tasks. *CogSci*.

Campbell, J., & Lee, M. D. (2006). The effect of feedback and financial reward on human performance solving 'secretary' problems. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *28*(28).

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical science*, *4*(3).

Gilbert, J. P., & Mosteller, F. (2006). Recognizing the maximum of a sequence [Reprinted from *Journal of the American Statistical Association (1966), 61*]. In *Selected Papers of Frederick Mosteller*. Springer.

Guan, M., Lee, M., & Silva, A. (2014). Threshold models of human decision making on optimal stopping problems in different environments. *Proceedings of the annual meeting of the cognitive science society*, *36*(36).

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4).

Guan, M., Lee, M. D., & Vandekerckhove, J. (2015). A hierarchical cognitive threshold model of human decision making on different length optimal stopping problems. *CogSci*.

Kocher, M. G., Praxmarer, M., & Sutter, M. (2020). Team decision-making. *Handbook of Labor, Human Resources and Population Economics*.

Kugler, T., Kausel, E. E., & Kocher, M. G. (2012). Are groups more rational than individuals? a review of interactive decision making in groups. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*(4).

Lee, M. D. (2006). A hierarchical bayesian model of human decision-making on an optimal stopping problem. *Cognitive science*, *30*(3).

Lee, M. D., & Courey, K. A. (2021). Modeling optimal stopping in changing environments: A case study in mate selection. *Computational Brain & Behavior*, *4*.

Lee, M. D., & Paradowski, M. J. (2007). Group decision-making on an optimal stopping problem. *The Journal of Problem Solving*, *1*(2).

Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. (2016). Probabilistic programming in python using pymc3. *PeerJ Computer Science*, *2*.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the" secretary problem". *Organizational behavior and human decision processes*, *69*(3).

Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, *13*(4).

Shupp, R. S., & Williams, A. W. (2008). Risk preference differentials of small groups and individuals. *The Economic Journal*, *118*(525).

Srivastava, N., Müller-Trede, J., & Vul, E. (2016). Modeling sampling duration in decisions from experience. *Proceedings of the Annual Meeting of the Cognitive Science Society*.

Thomas, B., Coon, J., Westfall, H. A., & Lee, M. D. (2021). Model-based wisdom of the crowd for sequential decision-making tasks. *Cognitive Science*, *45*(7).

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and computing*, *27*(5).