DC44 - Cohen's BEAST

Authors: Doron Cohen, Oded Cohen.

Contact information: Doronco30@gmail.com

MSD achieved by our model: 0.00584.

All 14 qualitative phenomena passed.

MSD for the competition set: 0.0088

First, we present a description of the baseline model, appearing also in Erev, Ert & Plonsky (2015). Then, we describe our unique contribution to the baseline's model predictions.

The baseline model, referred to as Best Estimate And Simulation Techniques (BEAST), assumes that Option A is strictly preferred over option B, after r trials, if and only if:

$$(1)[BEV_A(r)BEV_B(r)] + [ST_A(r)ST_B(r)] + e(r) > 0$$

Where BEVA(r) – BEVB(r) is the advantage of A over B based on the best estimation of the expected values, STA(r) – STB(r) is the advantage of A over B based on mental simulations, and e(r) is an error term[1].

In trivial choices, when one of the options dominates the other, e(r) = 0[2]. In all other cases e(r) is drawn from a normal distribution with a mean 0 and standard deviation σi (a property of agent i).

When the payoff distributions are known (the non-ambiguous problems in our study), the best estimations of the expected values are the actual objective ones. That is, BEVj(r) equals the expected value of option j, EVj (for all r). The simulation-based estimate of option j, STj(r), equals the average of κi (a property of i) outcomes that are each drawn (from option j's possible outcomes) in one mental simulation[3].

Each simulation uses one of four techniques. Simulation technique Unbiased implies random and unbiased draws, either from the options' described distributions or from the options' observed history of outcomes. Before obtaining feedback (decisions in trials 1 to 6) the draws are taken from the objective distributions using a luck-level procedure: The agent first draws a luck-level, a uniform number between zero and one. Then, for each prospect, the agent uses the same luck-level as a

---

[1] When the left-hand side of Inequality 1=zero, we assume random choice between the options.

[2] Dominance= either deterministic dominance or first-order stochastic dominance.

[3] E.g., consider an agent with κi = 3 who faces Problem 17 ("30" or "50, .5; -1") based on the following simulation results {30, 50}, {30, 50} and {30, -1} and the error term *e*(*r*) = -2. Equation 1 yields (30 – 24.5) + (90/3 – 99/3) – 2 = 0.5. Thus, the model implies an A choice.

percentile in the prospect's cumulative distribution function and draws the outcome that fits that percentile[4].

When the agents can rely on feedback (trials 7 to 25) they first sample one of the previous trials (all trials are equally likely to be sampled), and the drawn outcomes for both options are those observed in that trial.

The other three techniques are "biased": they can be described as a mental draw from distributions that differ from the objective distributions. The probability of choosing one of the biased techniques decreases when the participants receive feedback. Specifically, it equals:

$$(2) PBias(t) = \frac{\beta i}{(\beta i + 1 + t^{\theta i})}$$

Where $\beta i > 0$ captures the magnitude of the agent's initial tendency to use one of the biased techniques, t is the number of trials with feedback, and $\theta i > 0$ captures agent i's sensitivity to feedback[5].

The three biased techniques are each used with equal probability, PBias(t)/3. Simulation technique Uniform yields each of the possible outcomes with equal probability (see a related idea in Birnbaum, 2008) using the luck-level procedure described above (the draws are made from the uniform cumulative distribution function even after feedback is obtained).

Simulation-technique Contingent Pessimism is similar to the priority heuristic (Brandstätter et al., 2006); it depends on the sign of the best possible payoff (SignMax), and the ratio of the minimum payoffs (RatioMin). When SignMax > 0, and RatioMin ≤ γi ($0 < \gamma i < 1$ is a property of i), this simulation yields the worst possible payoffs for each option (MINA and MINB). This helps the model capture loss aversion and the certainty effect. When one of the two conditions is not met, the current simulation implies random choice among the possible payoffs (identically to technique Uniform). RatioMin is computed as:

$$(3) RatioMin = \begin{cases} 1 & if Min_A = Min_B \\ \frac{Min(|Min_A|, |Min_B|)}{Max(Min_A, Min_B)} & if Min_A \neq Min_B \quad and \quad sign(Min_A) = sign(Min_B) \\ 0 & otherwise \end{cases}$$

---

[4] i.e., the outcome drawn is the result of *F-1*(*x*), where *x* is the luck-level and *F* is the prospect's cumulative distribution function. For example, in Problem 2 ("3, .25; 0" or "4, .2; 0"), a luck level of .67 yields the draw {0, 0}, a luck level of 0.77 yields the draw {3, 0}, and a luck level of .87 yields the draw {3, 4}.

[5] E.g., assuming $\beta i = 3$, and $\vartheta i = .5$, the probability of using one of the biased techniques in each of the κi simulations is 3/(3+1) =.75 when *t* = 0 (trials 1 to 6), 3/(3+1+1) = .6 when *t* = 1 (Trial 7), and 3/(3+1+3.36) = .407 when *t* = 19 (Trial 25).

For example, RatioMin = 0 in Problem 9 ("1" or "100, .01; 0"), and 0.5 in Problem 10 ("2" or "101, .01; 1"). The contingencies capture two regularities. The sensitivity to SignMax implies less pessimism (less risk aversion) in the loss domain, hence the reflection effect. The second, RatioMin contingency, implies less pessimism when the minimal outcomes appear similar (have the same sign and are close in magnitudes). This implies that the addition of constant to all the payoffs, decreases risk aversion in the gain domain. In addition, it implies higher sensitivity to rare events in problems like Problem 10 and Problem 61 (large RatioMin), than in problems like Problem 9 and Problem 25 (small RatioMin).

Simulation technique Sign implies high sensitivity to the payoff sign. It is identical to technique Unbiased with one important exception: Positive drawn values are replaced by R, and negative outcomes are replaced by -R, where R is the payoff range (the difference between the best and worst possible payoffs in the current problem; e.g., 100 in Problem 9 and Problem 10)[6].

When the probabilities of the different outcomes are unknown (as in the problems with ambiguous Option B), they are initially estimated with a pessimistic bias (Gilboa & Schmeidler, 1989). The initial expected value of the ambiguous option is estimated as a weighted average of three terms: EVA, MINB, and UEVB, which is the estimated EV from Option B under the assumption that all the possible outcomes are equally likely. We assume the same weighting for EVA and UEVB, and capture the weighting of MINB with $0 \leq \varphi i \leq 1$: an ambiguity aversion trait of i. That is,

$$(4) BEV_B(0) = (1-_i)(UEV_B + EV_A)/2 + iMIN_B$$

For example, assuming φi = 0.05, BEVB(0) in Problem 22 (“10, .5; 0” or “10, p; 0”) equals .95(5+1)/2 + .05(0) = 2.85. In the no feedback trials (1 to 6) the probabilities of the m possible outcomes are estimated under the assumption that the subjective probability of the worst outcome SPMINB is higher than 1/m, and each of the other m−1 subjective probabilities equal (1- SPMINB)/(m-1). Specifically, SPMINB is computed as the value that minimizes the difference between BEVB(0) and the estimated expected value from Option B based on the subjective probabilities: SPMINB·MINB + (1- SPMINB)UBh, where UBh =(mUB- MINB)/(m-1) denotes the average of the best m−1 outcomes. This assumption implies that

$$(5) SP_{MIN_B} = \begin{cases} 0 & if BEV_B(0) > U_{Bh} \\ 1 & if BEV_B(0) < Min_B \\ \frac{U_{Bh} - BEV_B(0)}{U_{Bh} - Min_B} & otherwise \end{cases}$$

[6] E.g., in Problem 9 ("1" or "100, .01; 0"), all the positive outcomes are replaced by +100 (the value of *R*), and the 0 remains 0.

That is, in Problem 22 with φi = 0.05, SPMINB = (10 − 2.85)/(10 − 0) = 0.715. Each trial with feedback in the ambiguous problems moves BEVB(t) toward EVB. Specifically,

$$(6) BEV_B(t+1) = (1 - 1/T)BEV_B(t) + (1/T)O_B(r)$$

Where T is the expected number of trials with feedback (20 in the current setting) and OB(r) is the observed payoff generated from the ambiguous Option B at trial r[7]. The six properties of each agent are assumed to be drawn from uniform distributions between 0 and the model's parameters: σi ~ U(0, σ), κi ~ (1,2, 3, ..., κ), βi ~ U(0, β), θi ~ U(0, θ), γi ~ U(0, γ), and φi ~ U(0, φ). Namely the model has six free parameters: σ, κ, β, γ, φ, θ. Notice that only four of these parameters are needed to capture decisions under risk without feedback (the class of problems addressed by prospect theory). These parameters are σ, κ, β, and γ. The parameter φ captures attitude toward ambiguity, and θ abstracts the reaction to feedback. Best fit was obtained with the parameters σ = 7, κ = 3, β = 2.6, γ = .5, φ = .07, and θ = 1.

Our contribution

We observed that the BEASTs model predictions deviates systematically from the actual choices reported, on two different occasions;
1. When there are more than two possible outcomes in option B (i.e., *lotnum* > 1) or
2. The payoff distribution of option B in unknown (ambiguous problems, *amb*=1).

Our model adds an additional criterion that decides the direction of the deviations from the actual choices made by participants. This criterion is the difference between the lowest possible outcome of option A (*La*) and the expected value of the lottery option B (*Hb;* i.e. criterion= |*La-Hb*|)[8]. To correct for this deviation, we add a new parameter, *Diffbias*, dependant on the rules and criterion above.

We employ this two rules after the BEAST model, as described above, has produced its prediction of B choice rate for each of the problems. Then, for each problem, our model asks the following:

<u>If option B in the current problem has more than two possible outcomes (i.e. *lotnum>1)*</u>, the *Diffbias* parameter is added to each of the BEASTs' block predictions in the following manner: when |La-Hb| > 16, (-*Diffbias)* is subtracted[9]. When |La-Hb| <= 16, (+*Diffbias*) is added[10].

[7] E.g., in Problem 22 with *φ*i = 0.05, observing *OB*(6) = 0 implies that *BEV*B(1) = (1 − 1/20)·2.85 + (1/20)·0 = 2.707.
[8] E.g., in Problem 69, *La*=11 and *Hb*=31, thus |*La-Hb*|= |11-31| =20. In Problem 45, *La*=13 and *Hb*=13, so |*La-Hb*|=|0-0| = 0.
[9] E.g., for Problem 69, a choice rate of 0.8 predicted by BEAST will become 0.8–Diffbias.
[10] E.g., see Problem 45, where the final outcome will be the prediction made by BEAST+*Diffbias.*

Similarly, if the payoff distribution of option B is unknown (*amb=1*), then when |La-Hb| > 20, (-*Diffbias*) is added to BEASTs' final prediction for a given problem (*across all 5 blocks*), and when |La-Hb| <= 20, (+*Diffbias*) is added.

If a problem has both more than two possible outcomes, *and* is an ambiguous problem, *Diffbias* is added or subtracted (depending on the level of the criterion) **both times**. Thus, the process of adding or subtracting the parameter is serial and independent. For example, in Problem 69, *lotnum>1* and *amb=1*. Because for this problem |La-Hb|=|11-31|=20, first our first rule subtracts (-*Diffbias*) and then our second rule adds (+*Diffbias*).

The parameter *Diffbias* is a property of the agent, and is assumed to be drawn from a uniform distributions between 0 and: *Diffbias*i ~ U(0, *Diffbias*).

We obtained best fit with *Diffbias* = 0.07. After correcting the BEASTs' predictions, the MSD improved to 0.005, all 14 qualitative phenomena captured.

References

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review, 113*(2), 409-432. doi:10.1037/0033-295X.113.2.409

Erev, I., Ert, E., & Plonsky, O. From Anomalies to Forecasts: A Choice Prediction Competition for Decisions under Risk and Ambiguity.

Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, *18*(2), 141-153.