A Variation on the Paradox of Two Envelopes

Mikelis Bickis, Eric Neufeld

Department of Computer Science
University of Saskatchewan
Saskatoon, SK, S7K 5A9
bickis@snoopy.usask.ca, eric@cs.usask.ca

Abstract

The paradox of two envelopes, one containing twice as much money as the other, is one of several that address logical aspects of probabilistic reasoning. Many presentations of the paradox are resolved by understanding the need to specify the distribution of quantities in the envelopes. However, even if a simple distribution of numbers is known, a new set of reasoning problems arise. The puzzle as presented here shares features with other reasoning problems and suggests a direction for their resolution.

Introduction

The paradox of two envelopes is an old puzzle of probabilistic reasoning with a rich pedigree. We remark as a caveat that it is not a true paradox, but rather a puzzle whose statement is counterintuitive. It arises as a mathematical recreation, but has also studied by scholars in statistics (Christensen & Utts, 1992), philosophy (Rawling, 1997), and Artificial Intelligence (Neapolitan, 1990). It is one of many puzzles used to illustrate the logical foundations of probabilistic inference, as it shows the difference between uncertainty (in the sense of a probability distribution) and ignorance (lacking knowledge of even a distribution).

A typical presentation of the basic puzzle goes as follows: Ali and Baba have each been given an envelope of money. (Ali is male and Baba is female so we can refer to each with pronouns.) Both know one envelope has twice as much as the other and Ali has been offered an opportunity to switch envelopes with Baba. Assuming equal probabilities for receiving either envelope, Ali figures that he will improve his expected gain by switching, since

$$?* \ \boldsymbol{x}/2 + ?* \ 2\ \boldsymbol{x} = 5/4\ \boldsymbol{x}.$$

Following this reasoning, the first paradox that arises is that, having switched envelopes, Ali can follow the same reasoning again to increase the expected value of the contents of the envelope, and so on ***ad infinitum*** Hence, some call this the puzzle of the money pump.

The puzzle is discussed widely, but a short answer is that for the equation to hold for every possible value of $\boldsymbol{x}$, one requires a uniform probability distribution on an unbounded set. Since such a distribution does not exist, several papers pursue Bayesian approaches that postulate possible distributions. Under such models, the prior distributions of the contents of the two envelopes would be exchangeable, and hence both envelopes would have the same expected value. This solution cautions against na?ve use of the principle of indifference to generate probability distributions. We concur with this advice and do not discuss the classic puzzle further.

The present work in fact discusses solutions to a variation of the puzzle discussed by Bickis (1998). This variation differs from the classic version in that the distribution of the money in the two envelopes is fully specified in advance.

Although the new puzzle can be generalized, the following captures its essence. A number $\boldsymbol{x}$ is randomly chosen from the interval [0,100], say. The envelopes are filled with (arbitrary precision) cheques for $\boldsymbol{x}$ and $\boldsymbol{2x}$, then shuffled fairly and dealt to Ali and Baba who separately inspect the contents. Baba may then to ask Ali to switch envelopes. If Ali accepts, the envelopes must be switched. The first interesting variation on the puzzle arises simply because maximum values are known. If Ali sees a cheque greater than $100, switching can only result in a loss.

If Ali sees a cheque for less than $100, there is an argument that he may benefit from swapping envelopes. Let ***Largest*** be the event that Baba holds the largest cheque and let $\boldsymbol{f_1}$ and $\boldsymbol{f_2}$ be the density functions of $\boldsymbol{x}$ and $\boldsymbol{2x}$ respectively, Bayes theorem gives

$$P(Largest|A=a) = f_1(a) * P(Largest) \,/\, (f_1(a) * P(\sim Largest) + f_2(a) * P(Largest))$$

According to this, if $\boldsymbol{a <= 100}$ *?* $\boldsymbol{A=a}$, the probability that Baba holds the largest cheque is 2/3 and Ali would benefit from swapping, since $\boldsymbol{f_1}(\boldsymbol{a})$=1/100 and $\boldsymbol{f_2}(\boldsymbol{a})$=1/200.

Bickis (1998) points out that there is more to the decision than a simple (in the sense of sample space) counting argument. Logical information is available. Suppose Ali has $80 and Baba has $40 and Baba offers to switch. It does not follow that Ali should accept, for the fact that Baba makes the offer indicates that she has less than $100. Since she can only hold $160 or $40, her willingness to switch indicates that she has only $40 and that Ali should decline Baba's offer.

Suppose instead that Ali offers to switch. Regardless of probabilities, Baba will not accept if she is holding $160. If she does accept, Ali will know immediately that he has lost. That is, although a probabilistic argument exists that he might gain money, rationally he is guaranteed to lose. Therefore, he should not offer in the first place.

But then, by iterating the above arguments, one can deduce that regardless of the amount Ali gets, it is pointless for him to make an offer to switch, and that Baba should refuse any offer that is made. Ali would only make an offer if A <= 100, in which case Baba would refuse unless B <= 50, but in that case Ali would lose unless A <= 25. So that fact that Ali makes an offer tells Baba that A <= 25, so she would refuse unless B <= 12.50, but in that case, Ali would lose unless A <= 6.25 and so on.

It appears that the "logical" knowledge available contradicts the probabilistic analysis. The remainder of the paper provides a deeper analysis of the problem, suggests a solution and compares this problem to related problems in the literature.

A "backwards induction"

A few diagrams might be helpful. The cases for Ali holding either 160 or 80 are straightforward. The following diagram illustrates the possible lines of reasoning that may occur when Ali holds 40. He can reasons that Baba holds either 20 or 80.

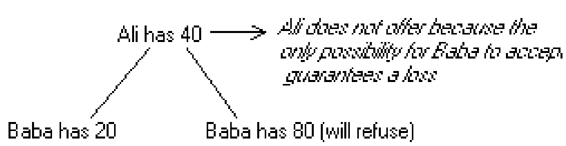

A rational Baba will hold onto 80, since she assumes Ali will not rationally offer to switch if he holds 160. Next suppose Ali has 20. Using the above diagrammatic

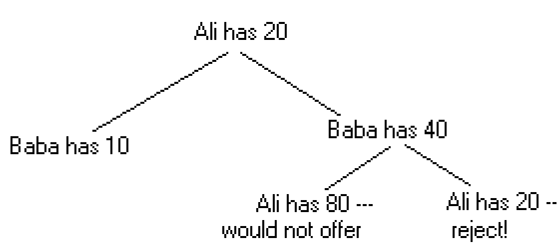

Once again, the only conditions where Baba might accept an offer are those where Ali is certain to lose, and so Ali will not rationally offer to switch at 20.

The arguments for Ali holding 160, 80 and 40 and 20 seem to rest on entirely rational (e.g., simple first-order) grounds. The next iteration is questionable:

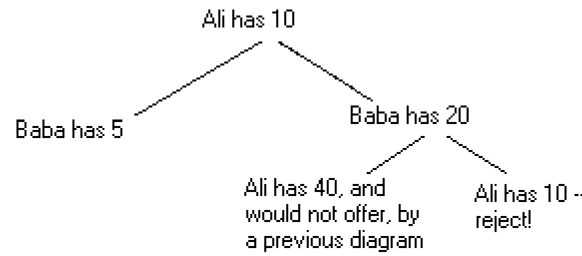

At this point the argument adopts a different modality, which is indicated by the note that "Ali would not offer ? ***by a previous diagram***". Ali is reasoning about how Baba would respond to an offer that Ali has not made, and, by this logic, will not make.

However, by iterating on this argument, the reader will see that Ali can ***always*** reason that he should never make an offer because Baba will only accept the offer when Ali is certain to lose the bargain. Stranger still, the argument is symmetric, so we could have two persons with differing (and possibly miniscule) amounts of money convinced that an accepted offer to switch guarantees a loss of money. This is distinct from the argument that there is no expected gain from switching. Yet the base case argument that a player with 160 should not offer to switch, and even the next case, appear indisputable. The problem is determining where, along this apparent "backwards induction" the reasoning fails.

Bickis (1998) gives an argument that it comes down to Baba having to decide how careless Ali's reasoning may be, and vice-versa. This provides a clue to the solution, although the problem is wider, since the argument that either can ***know*** in advance that the switch is hopeless for all sums seems flawed: consider the case of Ali and Baba both holding tiny cheques — how can it be ***known*** that the other will accept an offer of switching if the offerer is guaranteed to lose? Stranger still, it seems wrong that the very act of offering to switch guarantees a loss.

The flaw in the argument is revealed by attempting to represent the knowledge about the two-envelopes world in the formal language of first order logic. Suppose the initial knowledge is captured as follows:

1. ***Holds(A, X) -> (Holds(B, 2*X) or Holds(B, X/2).***
2. ***Holds(A, X) and X > 100 -> ~OfferToSwitch(A).***
3. ***Holds(A, X) and CantGain(A,X) -> ~OfferToSwitch(A).***
4. ***Holds(A, X) and Holds(B, Y) and Y > 100 -> Rejected(A, X).***
5. ***Rejected(A, X) ->CantGain(A,X).***
6. ***Holds(A, X) and Holds(B, Y) and Y < X -> CantGain(A, X).***

A complete and strict representation would include facts to support that each player can only hold a single quantity, the ranges of the quantities, and perhaps mostly importantly, explicit axioms about arithmetic. This encodes the basics of what we might call a minimally rational Ali and Baba, ignoring the argument of whether minimally rational humans have all of arithmetic at their disposal. Their shallow knowledge (as written) states that Baba holds either half or twice as much as Ali (1). (The axioms of arithmetic and knowledge about ranges of values allow a dual axiom for Baba.) If Ali holds more than 100, he won't offer to switch (2). It also lets him reason that if he can't possibly increase his wealth, he won't offer to switch (3). Baba will reject an offer if she is holding more than 100, that is, Ali's offer will be rejected (4). Finally, if Ali is rejected by Baba, he can't gain (5), and if Baba is holding less than Ali, he can't gain (6).

The first few proofs are straightforward. If Ali holds (say) 160, it follows from (1) that he will not offer to switch. If Ali holds 80, then Baba holds either 40 or 160 (4). Reasoning by cases, either Baba holds more than 100, or less than Ali. This implies Ali is either rejected or can't win, and therefore rationally should not to offer to swap.

Taking the argument further, suppose Ali holds 40. If Baba holds 80, she can reason using a contrapositive of (2) that Ali will not offer to switch if he holds 160 and therefore will reject any offer. Moreover, Ali can reason this much about Baba. Thus, the first three steps seem to define an easy induction.

Finally, suppose Ali holds 20. Then Ali can reason that Baba holds either 40 or 10 by (1). If 40, Ali then reasons that Baba reasons that a rational Ali with 80 would not have offered and that Ali holds 20 and would reject an offer, ***if one was made***. The other possibility is that Baba holds 10, in which case Ali's offer ***if made*** could either be rejected or accepted, with a worst case outcome of a loss.

Emphasis is added to illustrate that we are, at the level of discourse, reasoning about hypothetical true outcomes of predicates that will always be false as a consequence of the decision that the reasoning advises.

A feature of this line of reasoning is that it offers a technique of "backwards induction" that lets Ali conclude it is not worth making the offer to switch an envelope containing any amount. This bears some similarity to the puzzle of the unexpected hanging, where a prisoner is advised by the king that the prisoner will be executed one day at noon next week, and furthermore, the execution will come as a surprise. The prisoner reasons that the execution cannot be Saturday (last day of the week), since it would not be a surprise. However, this means that if the prisoner awakes Friday morning knowing a Saturday execution is impossible, then the prisoner must be hung Friday if the execution is to take place at all, in which case it would not be a surprise. The prisoner continues with another "backward induction", ultimately reasoning that there is no day the prisoner can be executed ***and*** be surprised. Hence, no execution.

This puzzle has its own pedigree and history of solutions that we do not review here, but see (Wischik, 1996), who references a solution that distinguishes between imagining actions, and implementing them. In the discussion above, Ali imagines making an offer to switch, then reasons about Baba's response. In his imagining of Baba's response, Ali assumes that Baba reasons that Ali would ***not*** make an offer in the event he was certain to lose, rules out the possibility that Ali holds a sum larger than Baba, and in fact, does ***not*** make the offer.

This notion of the difference between imagined and implemented actions suggests a resolution. Our first-order representation must distinguish between a capricious offer and a reasoned offer. Logic does not prohibit us from reasoning about imaginary worlds, but the essential features of the imaginary world must be the same as our own.

Thus, the previous representation becomes the following:

1. ***Holds(A, X) -> (Holds(B, 2*X) or Holds(B, X/2).***
2. ***Holds(A, X) and X > 100 -> ~LogicalOfferToSwitch(A).***
3. ***Holds(A, X) and CantWin(A,X) -> ~LogicalOfferToSwitch(A).***
4. ***Holds(A, X) and Holds(B, Y) and Y > 100 -> LogicalRejectOffer(A, X).***
5. ***RejectOffer(A, X) ->CantWin(A,X).***
6. ***Holds(A, X) and Holds(B, Y) and Y < X -> CantWin(A, X).***
7. ***OfferToSwitch(A) -> LogicalOfferToSwitch(A) or CapriciousOfferToSwitch(A)***
8. ***RejectOffer(A, X) -> LogicalRejectOffer(A, X) or CapriciousRejectOffer(A, X).***

The key feature of the new representation is that offers to switch are separated into logical and capricious offers. Symmetrically, rejections are divided the same way. (Additional axioms would ensure these are mutually exclusive and exhaustive, etc.) The conclusions of the backwards induction now agree with intuition. If Ali holds 160, the above tells him it would be illogical to offer to switch. However, this does not prevent Ali from making a capricious offer because of logical myopia or for sport. Baba, presented with an offer, before deciding whether to accept it or not, must also decide whether the offer is logical or capricious.

Going to the second case, suppose Ali has 80. If Ali makes a capricious offer, and if Baba has 160, she will logically reject the offer using (4), but this does not prevent her from capriciously accepting. If Baba has 40, she must first decide whether Ali is making a logical or a capricious offer before deciding to logically accept or reject. However, the scenario is similar to that one discussed earlier and it seems highly unlikely that Ali would make a capricious offer almost certainly knowing he will lose.

The reasoning changes slightly if we suppose Ali has 40. Suppose Ali makes a capricious offer. If Baba holds 80, then Baba reasons that Ali has either 160 or 40. It seems highly unlikely that Ali would make a capricious offer if he holds 160, and it seems highly reasonable for Baba to conclude that she should logically reject the offer. But recall that this is Ali reasoning about Baba reasoning about whether Ali has made a reasoned offer. If we pursue the argument for smaller values, it is difficult to characterize the same events as highly likely or highly unlikely.

Thus, the special case involving certainty (when one player holds 160) has devolved to a case where it is highly unlikely (but not impossible) that a player would make an offer to switch given a certainty of losing. By analogy: two reasonably good (perfectly rational) tic-tac-toe players can ensure that every game ends in a tie, regardless of the opening move. However, it seems believable that one of two intelligent but inexperienced players might make a losing move in response to an opening, even if that same player would not deliberately make a last move that forced his opponent's win. The tic-tac-toe analogy ends there because that game is finite.

Relationship to other work

Two features of the new puzzle bear some relationships to other work.

Glenn Shafer (1985) discusses the idea of ***communication conventions*** in the context of determining probabilities in the puzzle of the two aces. In that puzzle, a deck of cards consists of an ace of hearts, an ace of spades, a two of hearts and a two of spades. The deck is shuffled and two cards are dealt to a player ***A.*** Player ***B*** asks ***A*** whether ***A*** holds an ace. Player ***A*** answers, "yes, in fact I hold the ace of spades". Player ***B*** then computes the probability that ***A*** holds the other ace. Shafer's discussion revolves around whether the correct answer is determined by computing

P(holds(A, aceOfHearts) | holds(A, aceOfSpades))

or by computing

P(holds(A,bothAces) | holds(A,oneAce)).

Had player ***A*** narrowly answered ***B's*** question by simply replying "yes", it would only be possible to use the second probability. However, ***A*** has ventured some additional information, and the paper revolves whether ***B*** can use the additional information in the solution. Shafer conclude that ***B*** can not, since the probabilistic communication implies a simple yes/no answer, and ***A*** might throw out information to mislead ***B*** as much as help ***B.*** Thus, it is different to discover find information you are looking for, than to chance upon it. In terms of an objective interpretation of probability, it would be difficult to fully specify the sample space in advance if one must give an objective account of the possible intentions of all players. In this new puzzle of two envelopes, the reasoning of Ali and Baba also rests on the way information was obtained.

There also seems to be a link between the backwards induction of the unexpected hanging and the backwards induction of our puzzle. One potential flaw in the reasoning is that the prisoner is reasoning that the judge has advised that the prisoner will die ***and*** will be surprised. The prisoner reasons that since there cannot be a surprise, there cannot be an execution, contrary to the advice of the judge. But the judge has advised ***both,*** yet the prisoner does draw the less convenient conclusion that since there must be an execution, there cannot be a surprise.

Shapiro (1998) notices a problem with the induction, as does (Wischik, 1996). Both propose parallel solutions to the problem that seem to work, except for the case of the execution occurring Saturday. (The victim will still be surprised, but for other reasons --- the executioner must have lied, etc.) Shapiro's solution is neat. However, it is not clear whether the puzzle cannot be posed again, assuming that the rational prisoner has Shapiro's parallel reasoning formalism available.

Putting that aside, the meaning of the puzzle is resolved by casting it into a formal language. The following is a bit freewheeling, but captures the main features of the argument. First, the sentences establish that the prisoner is alive now, but will hang on some day in the upcoming week:

1. ***Hang (1) or Hang(6) or ? or Hang(7)***
2. ***Alive(0)***
3. ***Alive(N) -> Alive(N-1)***
4. ***Hang(N) <-> ~Alive(N)***
5. ***Hang(N) ->Surprise(N).***

This captures some of the salient features of the effect of hanging. Implicitly, by contraposition, (3) also states that if you are not alive at ***N-1*** then you are also not alive at ***N.*** This would seem to be straightforward enough. However, this set of axioms asserts, perhaps stupidly, that a hanging is a surprise, albeit a nasty one, whether one is expecting it or not.

A better first-order definition of "surprise" is a worthy puzzle in itself. In the above, a surprise could be when you discover information you simply did not know (equivalently, you determine an atom, for example, ***Hang(X)***, is true). Alternately, and more appropriately for this puzzle, a surprise might be defined as discovering that the truth of a fact in the world is contrary to a proof in the language ***L*** consisting of the above axioms plus the first-order logic.

Kyburg (1984) says that every experiment is at once a test of a hypothesis ***and*** of a measurement. If the measurement obtained by the experiment contradicts the theory, when do we toss out the theory and when do we toss out the data? (Pople (1982) states that physicians discard data in the process of differential diagnosis once a certain amount of effort has been invested in a hypothesis. This puzzle is not different. On one hand, the puzzle defines a ***prima facie*** sensible theory of hanging and surprises.)

There are many avenues to pursue in trying to replace Axiom 5. One approach that follows our solution to the Ali-Baba paradox is to define two kinds of surprise, first-order surprise (***surprise1)*** and second-order surprise ***(surprise2)***. A first-order surprise occurs on day ***N*** if the prisoner cannot use axioms 1 to 4 to deduce ***whether*** the hanging will occur on day ***N,*** and that the hanging ***does*** occur that day. The following first-order formulation is clumsy, but serves. To handle this,

5'.HangPossible(M != N)
& HangPossible(N) & Hang(N) -> Surprise1(N).

The first predicate is shorthand notation meaning that if hanging is possible on some day other than day ***N***, then it is possible for a ***surprise1*** to occur. For now, we will simply assert that hanging is not possible on the seventh day:

5.1 ~HangPossible(7).

Finally, we need to add some notion of time. We add

5.2 OnDay(N) -> ~HangPossible(N-1) & ~Hang(N-1).
This just says that if day N has happened, hanging was not possible the previous day. We also rewrite 5' and 5.2 as

5.3 OnDay(N) & HangPossible(M != N)
& HangPossible(N) & Hang(N) -> Surprise1(N).
5.4 OnDay(7) ->~HangPossible(7).

We do not complete the biconditional in 5.3. This makes it possible to state that if the prisoner is not hung on day ***N***, then hanging is not possible on day ***N***. This is clumsy, but spares us the burden of adding time to the ontology. On the other hand, if hanging is not possible on day ***N (~HangPossible(N)),*** we still wish ***Hang(N)*** to be possible.

The predicates ***Hang()*** and ***HangPossible()*** decouple the prisoner's reasoning ability from actual reality. We could then define a second-order surprise as follows:

5.3 OnDay(N) & ~***HangPossible(N) & Hang(N) -> Surprise2(N).***

However, it is fairly straightforward to show that ***HangPossible(N)*** is false for all ***N*** and thus, a first-order surprise is not possible. However, this does not preclude a second order surprise, which we can simplify, using the knowledge that ~***HangPossible(N)*** is always false, to the almost trivial

5.3 Hang(N) -> Surprise2(N).

This definition, surprisingly, is identical to our first definition of ***Surprise*** in the "stupid" axioms, that a hanging is always a surprise. It seems to make sense in a world where we have deduced that hanging is not possible.

Conclusions and Future Research

The puzzle of the two envelopes continues to interest many scholars. The new version of the puzzle presented here rests on an apparently reasonable reverse induction on the quantity of money in the envelope. The flaw in the argument is not with the induction, but on the fact that the teller of the puzzle is identifying a reasoned offer to switch envelopes with a hypothetical arbitrary offer. The subsequent solution of clearly distinguishing the kinds of offer in a formal language appears to have some application in disambiguating other puzzles.

Acknowledgements

The research of both authors is supported by grants from the Natural Sciences and Engineering Research Council of Canada. The authors thank students and colleagues for discussions and the referees for insightful comments.

References

Bickis, Mikelis. 1998. The Real Paradox of Ali and Baba. ***SSC Liaison***, **12** (2) June 1998, 20-25

Christensen, Ronald., and Jessica Utts. 1992. Bayesian resolution of the "exchange paradox". ***The American Statistician***, **46** (4) 274-278.

Kyburg, Henry E., Jr. 1984. ***Theory and Measurement.*** Cambridge University Press

Neaopolitan, Richard. 1990. ***Probabilistic Reasoning in Expert Systems.*** John Wiley, New York

Pople, Harry E., Jr. 1982. Heuristic methods for Imposing Structure on Ill-structured Problems. The Structuring of Medical Diagnostics. In ***Artificial Intelligence in Medicine,*** ed. Peter Szolovits, 119-190

Rawlins, Piers. 1997. Perspectives on a Pair of Envelopes. ***Theory and Decision*** **43** 253-277

Shafer, Glenn. 1985. Conditional Probability. ***International Statistics Review*** **53** (3) 261-277

Shapiro, Stuart. 1998. A procedural Solution to the Unexpected Hanging and Sorites Paradoxes. ***Mind 107*** **(428)** 751-761

Wischik, Lucian. 1996. ***The paradox of the surprise examination.***
http://www.wischik.com/lu/philosophy/surprise-exam.html