

City University of New York (CUNY)

CUNY Academic Works

Publications and Research

CUNY Graduate Center

2022

Logic, Co-ordination and the Envelope of our Beliefs

Rohit J. Parikh
CUNY Graduate Center

[How does access to this work benefit you? Let us know!](#)

More information about this work at: https://academicworks.cuny.edu/gc_pubs/727

Discover additional works at: <https://academicworks.cuny.edu>

This work is made publicly available by the City University of New York (CUNY).
Contact: AcademicWorks@cuny.edu

Logic, Co-ordination and the Envelope of our Beliefs

Rohit Parikh

CUNY

November 1, 2022

Abstract

Each of us has a story which we can think of as a set of beliefs, hopefully consistent. We make our decisions in view of our beliefs which may be probabilistic, in the general case, but simple yes or no as in this paper. Our beliefs are our envelope just as the shell of a tortoise is its envelope.

Decision theory - or single agent game theory tells us when to make the best choice in a game of us against nature. But nature has no desire to further or frustrate our efforts. Nature is mysterious but not malign.

Things change when there are other agents involved. Then the best thing for us to do will depend on *what they do*. And they will think the same. And we predict their actions in terms of what we think their beliefs are.

But how do we *coordinate* with others whose beliefs are different? This paper addresses the issue of *working together despite different beliefs*.

1 Introduction

Each of us carries with us an envelope of beliefs, just as a tortoise carries its shell. Working with others, even when we have compatible ends, requires reconciling our respective shells.

“How different is she from me?” That will depend on whether we are choosing a movie to go to or whether we are choosing whether to vote Democrat or Republican. This insight makes it possible to define a *logical distance* between two people, in areas in which they can collaborate, and other areas where they will see each other as opponents if not as enemies. The resulting metric space, or collection of metric spaces can be a subject of mathematical investigation.

For an example, consider people who refuse to be vaccinated because they feel that vaccination is dangerous. We may reproach them and try to convince them of the error of their beliefs. But again, if their belief is sincere, we would be more charitable. These points, about the connection between knowledge and obligation are spelled out in my paper [6] with Pacuit and Cogan. But it may be that we will consider the action of being vaccinated acceptable or even good and they will not. In that case their action is reasonable *given their beliefs*. Their choice of a different action is explainable in terms of their different beliefs. Here is a relevant quote from Robert van Rooij [8]:

In decisions under strict uncertainty and ordinal preferences, the decision problem can be modeled by a triple like (T, E, \geq) , where (i) T is the set of states that the agent thinks are possible, representing her beliefs, (ii) E is the set of alternative actions that she considers, and (iii) \geq is a partial order on state-action pairs, i.e. $\geq \subseteq T \times E$, which represents her preferences

Following Savage (1954), I will take the actions in set E to be primitives. When our agent can represent her preferences by a cardinal utility function, the ordering \geq is replaced by U , a utility function from state-action pairs to real numbers. When the agent can quantify her uncertainty, the problem is standardly called a decision under risk. In these cases a decision problem contains an additional probability function P which assigns to states their (subjective) probabilities”

van Rooij 2003

2 Ordering and preference

In the following, the variable t ranges over states of nature. The variable e ranges over possible actions we can take. We assume that there is a preference over outcomes of particular actions in particular states.

We recall that $(e, t) \leq (e', t)$ means that in state t , action e' is at least as good as action e .

Definition 2.1 *Given two actions e, e' and a set A of states we define $e \leq_A e'$ as $\forall t \in A, (e, t) \leq (e', t)$
(If A is a formula, we will identify it with the set of states which satisfy A).*

So $e \leq_A e'$ means, “given A , e' is at least as good as e ”

It is easy to see that if $B \rightarrow A$ and $e \leq_A e'$ then $e \leq_B e'$. For the states satisfying B are a subset of states satisfying A .

We will say that e is **maximal** given A iff $\forall e', e' \leq_A e$.

It follows that if e is maximal given A and $B \rightarrow A$, then e is also maximal given B .

Clearly, when we choose an action, we prefer to choose a maximal action, and learning more makes it easier to find a maximal action. However, maximal actions need not always exist. They are rather like dominant strategies which do not always exist.

2.1 Order versus utility - roulette lotteries

Suppose that Sona prefers chocolate ice cream to vanilla and vanilla to strawberry. Then we know that her ordering is $C > V > S$. But we do not know if her utilities are 10, 9, 1 in which case vanilla is almost as good as chocolate, or they are 10, 2, 1 in which case vanilla is not much better than strawberry. But order by itself does not allow us to define utilities. However, if we do have utilities, then we can define $(e, t) \leq (e', t')$ by $U(e, t) \leq U(e', t')$.

To get at the utilities, we can offer her a choice. a) she can just pick vanilla.

Or b) we toss a coin and if it lands heads she gets chocolate and if it lands tails she gets strawberry.

If she picks a) then her utilities are 10,9,1 (or similar). If she picks b) then her utilities are closer to 10,2,1. For details see Savage [9].¹

2.2 Horse lotteries

Suppose that there are three horses running. Baroda, Medina and Donegan.

Again we could offer her a choice. c) she gets vanilla or d) if Baroda wins, she gets chocolate, if Medina wins she gets vanilla and if Donegan wins she gets strawberry.

Whether she picks c) or d) will depend not only on her preferences in ice creams but also on her opinions of the horses. See Anscombe and Aumann [1].

In the following we will focus on roulette lotteries, and the probabilities will be *objective*.

2.3 Utility and Decision

Suppose that the ordering is created by a utility which is a number. So $(e, t) \geq (e', t')$ iff $U(e, t) \geq U(e', t')$ Then we can define the worst outcome (**security**) of an action e as $S(e) = \min_t U(e, t)$ and we can define the action with the best security to be e^* such that $S(e^*) = \max_e(S(e))$. There always exists an action which has the best security although it may not be unique.

More generally, $S_A(e) = \min_t U(e, t) : t \in A$. That defines the security of action e when A is given.

If B implies A then $S_B(e)$ is at least as great as $S_A(e)$

¹If there are utilities then it also becomes possible to compare $U(e, t)$ with $U(e', t')$ where t, t' are distinct. Now we can ask people to judge even with $t \neq t'$. We could say to an agent, “if the state were t' would you prefer e or e' ?” E.g. “if Zubin Mehta was the conductor rather than Seiji Ozawa, would you prefer the concert or a movie?” But such tests are not available with animals. We humans are Popperian and make choices in hypothetical situations. It is not clear if animals do.

In the following matrix, e_1, e_2 have security level zero and e_3 has security level 2. So e_3 is safest.

	e_1	e_2	e_3
t_1	4	1	3
t_2	6	0	2
t_3	0	5	4

However if we find out that t_3 is not in the running, e_1 with the resulting security level of 4 becomes the winner. Indeed the security level always rises when we learn something. For the minimum over a smaller set is always as great as the minimum over a larger set.

Clearly two agents with the same utility but different information can differ on which is the safest action. Suppose that the actual t is t_1 . One agent knows that it is not t_2 and another that it is not t_3 . Then hen they will disagree on the best action even if they have the same preferences and true beliefs. If we know that it is not t_3 then e_1 with safety of 4 is best. If all we know is that it is not t_2 then e_3 with security of 3 is best.

2.4 Probability

If in addition to the utility function U on such pairs we have a probability p defined on states (and hence on sets of states), then we can define an *expected value* for an action.

$$\bar{U}_i(X, a) = \frac{\sum_{w \in X} u_i(w, a) \cdot p(w)}{p(X)}$$

However, the expected value does not necessarily increase when we know more. When we know more, we might eliminate some good options and the expected value could fall. In the example above, the expected value of e_3 is 3. But if we find out that the state is not t_3 then the expected value falls to 2.5.

2.5 Is knowledge always beneficial?

Knowing more is not necessarily an advantage if other players are involved and they know that you know more.

Example: two individuals i and j are to predict whether the top card of a shuffled deck is black or red. The predictions are made in sequence: i guesses first, and then j (after hearing i 's guess). If they make the same prediction, each wins 1 euro, whether their prediction turns out correct or not. If they disagree, the one who is correct wins 3 euros, and the other nothing. Since i has no basis for preferring one or the other color, he will choose at random. Then j will choose the opposite color, so each will have an expected gain of 1.50 euros. Now suppose someone offers, without charge, an arrangement whereby the first player i is told the actual color of the top card before he makes his prediction. Evidently, the parties would be unanimous in rejecting that arrangement, even i as the beneficiary. For, with that information, individual i would choose the correct color, j would then make the same choice, and they would each gain only 1 euro rather than an expectation of 1.50 euros. (See Kamien et al [3])

Suppose however that i is told the color of the top card, and j is *not aware*. Then i will choose the right color, j will choose the opposite color and their payoffs will be 3 and 0.

So i loses by knowing the correct color but only if j knows that i knows. (See Neyman 1991, [5])

Question for the reader: Suppose that i knows the right color and knows that j knows that i knows. Is there a mixed strategy which i can use which has an expected payoff greater than 1?

3 Decisions and Beliefs

Whether a particular action is good or bad depends on what we believe. If we believe that a glass of wine is poisonous then we will not drink it. When more than one agent is involved and they are considering a pair of actions, they will want the pair to be good for both of them. Otherwise one of them will opt out. Also, if several pairs of actions are good then they will want

a pair which is *Pareto optimal* (in a sense to be defined). They may also want the benefit to be as good for each as possible and following Nash we will define the notion of *Nash optimality*.

If agents have different beliefs (but the same utilities) then a pair of actions considered good by one may be considered bad by the other. Or simply not good enough. Some form of dialogue will then become necessary so as to resolve the difference in beliefs. The beliefs do not need to coincide for us to agree on a pair of actions. If we are choosing among several movies, then our difference in views between baseball and cricket, no matter how large, may not be an obstacle.

Beliefs are expressed in language, so suppose that you and I share some language L .²

X is the set of my beliefs in L . Y is the set of **your** beliefs in L . $Z = X \cap Y$ is our common ground. When discussing with you I am entitled to use members of Z and also, members of X which I have convinced you of. I may also convince you to remove members of Y whose falsity I have convinced you of. I may **not** use elements P in X such that $\neg P$ is in Y unless I have convinced you of P (I leave aside for now the issue of whether membership in X or Y is common knowledge).

Now, as we shall see, whether a particular action is good or bad, or whether a particular pair of actions is good or bad for one or both of us, will depend on our beliefs. We will investigate the connection between beliefs, and judgments about actions.

Suppose that one of us (say you) does an act a which has utility functions for both of us, $f(e, t)$ for me, $g(e, t)$ for you. The action is **acceptable** if the expected values of $f(e)$, $g(e)$ are both positive.³ It is **strongly acceptable** if it is acceptable and (as with Nash bargaining) $f(e) \times g(e)$ is maximized relative to other possible actions of yours - including no action).

However, these notions are relative to beliefs. Suppose my total beliefs amount to some set of beliefs X . Yours to Y . If $\mathcal{B} = \{w | w \text{ is a world}$

²Given a belief A as a sentence we can identify it with the set of all states which make A true.

³Let us put aside the case where $g(e)$ is negative but you **altruistically** do e because you believe that $f(e)$ is positive.

and satisfies all elements of X } then as far as I am concerned, states in \mathcal{B} are all that exist. I should compute $f(a) \times g(a)$ over \mathcal{B} and not over [all possible worlds](#).. Clearly the calculation for me will depend on whether X or Y or Z are common knowledge. If I know Z but not Y , I can propose an action whose safety value given Z is positive for you, and then it will also be positive given $X \cup Y$. (Every world w satisfying $X \cup Y$ satisfies $Z = (X \cap Y)$).

But it is possible that an action a is has positive utility for both of us over the worlds which satisfy X but not over the set of worlds which satisfy Y . Then if I do a , you may suspect me of not being cooperative.

Similarly if you assert Q such that Q makes your action a acceptable but I do not believe that you believe Q then I will suspect you of hypocrisy. If the salesman tells you that a certain car is good value, but does not himself believe that, then surely he is guilty of hypocrisy.

I considered a *single action* done by just one of us. But perhaps there is a *pair of actions*, a, b which we can jointly perform. For instance in the Bach-Stravinsky game, either both listen to Bach or both to Stravinsky. The case where one listens to Bach and the other to Stravinsky has zero payoffs for both and need not be considered *if mutual consultation is possible*.

	Bach	Stravinsky
Bach	2, 1	0, 0
Stravinsky	0, 0	1, 2

The framework: We assume there are two players A and B (or Ann and Bob). There is a finite space W of possible states and a finite set A of possible actions. p is the probability on W . There are two utility functions u_1 and u_2 such that $u_i(w, a, b)$ is the utility to i if the pair of actions a, b is performed and the actual state is w . A, B are negotiating which pair of actions to perform depending on the expected consequences to them.

Let X be a subset of W . Then $\bar{U}_i(X, a, b)$ is the expected value of $u_i(a, b)$ over worlds in X . Specifically

$$\bar{U}_i(X, a, b) = \frac{\sum_{w \in X} u_i(w, a, b) \cdot p(w)}{p(X)}$$

We will be interested in those cases where A is a formula in our language L and X_A is the set of worlds (or states) satisfying A . Abusing notation we will write $\bar{U}_i(A, a, b)$ replacing X in our earlier notation by X_A .

Let $Con(a, b, A) = (\bar{U}_1(A, a, b), \bar{U}_2(A, a, b))$ be the pair of expected consequences for Ann and Bob if the pair of actions a, b are performed and A is true.

Example: Ann is suggesting going to the movie *A Day at the Races* at the movie theatre Roxy. The usual ticket price is \$15 per person and cost conscious Bob says, “it is too expensive”. Ann responds with A , “today is Thursday and they have a 50% discount on Thursdays”. Her remark reduces their total expected cost from \$30 to \$15.

Suppose Bob responds with B , “you are right about the rule, but today is actually a Friday and there is no discount.” So they have a disagreement about which formula A or B is true and hence of the expected cost.

Now their dispute can be resolved by simply consulting the calendar. Suppose that Ann finds the calendar and shows Bob that it is not Friday but Thursday. Then Bob replaces his formula B by A and they agree on the expected cost. Note that the conflict was resolved because Bob *already* had the belief “the calendar is correct.”

Initially they disagreed on the cost of the movie because they disagreed on which day of the week it was. But then their beliefs converged as did their estimate of the cost of the movie.

Example 2: Bob is buying a car and the salesman offers a model which Bob likes, for \$35,000. Bob is uneasy about the price and then the salesman says C , “this model is very reliable and will have very low repair costs. In fact for only \$500 we will cover all repair costs for the next ten years.”

Suppose the **usual cost** of repairs is \$3,000. Then Bob decides to buy the car as the sentence uttered by the salesman reduces the expected cost over ten years from \$38,000 to \$35,500. The salesman (if he is believed) makes the car a good bargain.

Suppose now that Ann and Bob have pairs of actions $(a_1, b_1), \dots, (a_m, b_m)$

available to them and are discussing which pair to perform.⁴ The expected benefits for them can be written $(c_1, c'_1), \dots, (c_m, c'_m)$ where c_i is the benefit to Ann if the pair (a_i, b_i) is performed and c'_i is the benefit to Bob. We will assume that the c_i, c'_i are all positive from their individual points of view, for otherwise they will not bother to act.

Single Belief case: We first consider the case where Ann and Bob have the same beliefs. Let X be the conjunction of their common beliefs. Then we might as well think that W , the set of worlds is simply the set of worlds which satisfy their common beliefs. We now define the notions of [Pareto-optimality](#) and [Nash-optimality](#).

Definition 3.1 *A pair of actions (a, b) is [Pareto optimal](#) if there is no other pair of actions (c, d) which yields consequences for each of Ann and Bob which are just as good as (a, b) for both Ann and Bob and strictly better for one of them.*

A pair of actions (a, b) is [Nash optimal](#) if there is no other pair of actions (c, d) which has a greater product of payoffs.

⁵ It is clear that a Pareto optimal pair need not be Nash optimal. For instance if \$100 are divided between Ann and Bob, Ann gets \$99 and Bob gets \$1 then this division is Pareto optimal but not Nash optimal. It is [Pareto optimal](#) because Ann would suffer with another division, say \$98 for her and \$2 for Bob. It is not Nash-optimal because each getting \$50 would give the higher product of 2,500. (2500 dollar square?)

Proposition 3.2 *Nash optimal does imply Pareto optimal.*

For suppose that utilities u, u' are Nash optimal but not Pareto optimal.

⁴We do not assume that (a_1, b_2) is necessarily also a pair as *that combination* might make no sense for them.

⁵See [4] for a discussion of Nash bargaining. It is easy to see that what matters is the expected value of the product rather than the product of the expected values. Thus if there are two possibilities, (99,1) and (1,99) then the product of the expected values is about 2,500. But the expected value of the product is only 99. And it is the latter which matters for fairness.

Then there exist (say) v, v' such that $u \leq v, u' < v'$. But then $u \times v < u' \times v'$ and (u, v) was not Nash optimal.

Nash used a convexity argument to prove the existence and uniqueness of a Nash optimal point. We on the other hand are using the finiteness of the set of pairs of actions. There could of course be [ties](#) which Nash did not have to worry about.

I do not know if the Nash bargain has been used in the literature to define fairness.

However, they may have different beliefs, and some c_i evaluated by one set of beliefs might be positive while it is negative by another set of beliefs.

Many belief case: This case will arise when the two parties do not have the same beliefs (but the same utilities). Let us suppose that Ann and Bob each have their own sets of beliefs, X for Ann and Y for Bob. In that case the notions of Pareto optimality and Nash optimality will come in two flavors. Conflicts will arise when one party tries to convince the other party of the wrongness of their beliefs.

One way one can convince the other is when the beliefs of one or other party are not consistent as a set. Suppose that Ann's set X is not consistent. When a question comes up to her, she uses a [consistent subset](#) $X' \subseteq X$ to resolve it.

Suppose now that Ann has a belief A and Bob has a belief B which is incompatible with A . Bob can then point to a set $X' \subseteq X$ such that X' implies $\neg A$. Ann will then be compelled to withdraw A and may agree to go along with B .

Bob's acceptance of the calendar as an authority was such an example. (It was Bob who was inconsistent in that example). Such themes have been developed by Dung, Amgoud, and others.

A conflict in beliefs is not necessary for different evaluations of the same action. If Bob's set of beliefs Y is a superset of Ann's set X then Bob accepts every belief that Ann has. And yet the expected value over the worlds which satisfy X and those over the worlds which satisfy Y may end up having different signs.

Here is an example. There are three states of the world, equally likely

- cold
- warm and sunny
- warm and rainy

The utilities of a walk (for both Ann and Bob) are, respectively, -2, +5, -2. Ann knows that it is warm. Bob knows that it is warm and rainy. All of Ann's beliefs are Bob's beliefs as well. The expected value of a walk as computed by Ann is +1.5, $((5-2)/2)$. As computed by Bob, it is -2.

Distance between beliefs Suppose we are given actions (a_1, \dots, a_m) . Ann believes A and Bob believes B . But they have the same utilities.

Let c_1, \dots, c_m , be the payoffs respectively of actions (a_1, \dots, a_m) calculated over the set of worlds satisfying A . d_1, \dots, d_m the payoffs calculated over the set of worlds satisfying B . Then we define the *distance* $d(A, B)$ between A, B relative to these actions as $\max(|c_1 - d_1|, \dots, |c_m - d_m|)$.

The distance is 0 if A and B are logically equivalent. It is symmetric in A, B so that $d(A, B) = d(B, A)$ and satisfies the triangle inequality.

To see the last, consider three formulas A, B, C . Now $d(A, C)$ is the maximum of certain differences, so let $d(A, C)$ equal $|c_k - f_k|$ where f_1, \dots, f_m , are the payoffs calculated over the set of worlds satisfying C . Now

$$d(A, C) = |c_k - f_k| \leq |c_k - d_k| + |d_k - f_k| \leq d(A, B) + d(B, C)$$

Thus we have a metric space defined by (L, d) where L is the language of communication (and we identify logically equivalent formulas where the distance will be zero). and d is the distance just defined.

Note that the distance depends not only on A, B but also on the sets of action pairs being considered. Let U be a set of action pairs and V be a subset of U . Let us denote the distances relative to U and V by $d_U(A, B)$ and $d_V(A, B)$.

Then given that $V \subseteq U$, $d_V(A, B) \leq d_U(A, B)$. The fewer the number of issues, the less the disagreement. A Ukrainian and a Russian can both agree

to have their coffee black but perhaps not on some other things. Also, One computational model suggests that the languages Russian and Ukrainian share about 55% of their vocabulary. And yet as we all know there are sharp differences about other matters.

Clearly if Ann and Bob have the same beliefs and same utilities then their values of payoffs of actions a_1, \dots, a_m (or pairs of actions $(a_1, b_1), \dots, (a_m, b_m)$) will be the same. They will agree on which actions are acceptable and which are Nash-optimal. But even if X are the beliefs of Ann and Y are the beliefs of Bob and $d(X, Y)$ is *small* then they may **still agree** on which pairs of actions are acceptable and which are Nash optimal.

Again, it may also happen that they agree on certain issues and disagree on some others. In that case it may happen that they will agree on which movie to go to and disagree on which school to send their children to. Such differences may arise as a result of a difference in utilities, but in this work we are focusing on a difference in beliefs.

Conclusion We have offered a framework in which beliefs affect the values of actions and different beliefs by two agents may result in a disagreement over which actions are good. We define the notions of Pareto optimal and Nash optimal. These notions are unambiguous when the utilities and beliefs are shared. They may be interpreted differently when the agents have different beliefs.. We can also notice that when Ann and Bob have different sets of beliefs they may agree somewhat on the utility value of one action pair (a, b) and disagree on the value of another pair (c, d) .

Issues having to do with knowledge and obligation have been addressed in [2,6]. Stalnaker [10] addresses the question of when we should believe another person. However our current treatment has some new ideas.

This work is quite preliminary and a more detailed treatment with more mathematics will emerge eventually.

Acknowledgements: we thank Neil Hwang, Ali Khan, Paul Pedersen and a referee for useful comments.

Selected references

1. Anscombe, Francis J., and Robert J. Aumann. "A definition of subjective probability." *Annals of mathematical statistics* 34.1 (1963): 199-205.
2. Grossi, Davide, et al. "How Knowledge Triggers Obligation." *International Workshop on Logic, Rationality and Interaction*. Springer, Cham, 2021.
3. Kamien, Morton I., Yair Tauman, and Shmuel Zamir. "On the value of information in a strategic conflict." *Games and Economic Behavior* 2.2 (1990): 129-153.
4. Nash Jr, John F. "The bargaining problem." *Econometrica: Journal of the econometric society* (1950): 155-162.
5. Neyman, Abraham. "The positive value of information." *Games and Economic Behavior* 3.3 (1991): 350-355.
6. Pacuit, Eric, Rohit Parikh, and Eva Cogan. "The logic of knowledge based obligation." *Synthese* 149.2 (2006): 311-341.r
7. Parikh, Rohit, Cagil Tademir, and Andreas Witzel. "The power of knowledge in games." *International Game Theory Review* 15.04 (2013): 1340030.
8. Van Rooy, Robert. "Quality and quantity of information exchange." *Journal of Logic, Language and Information* 12.4 (2003): 423-451.
9. Savage, Leonard J. *The foundations of statistics*. Courier Corporation, 1972.
10. Stalnaker, Robert. "Saying and meaning, cheap talk and credibility." *Game theory and pragmatics*. Palgrave Macmillan, London, 2006. 83-100.