COSC410 Lecture 4
Nonmonotonic Logic

Willem Labuschagne
University of Otago

2016

Introduction

Let us return to the Traffic System. The object language is $L_{\{p,q\}}$ where p stands for "The light is red" and q for "The oncoming car stops". The states of the system are given by $S = W_{\{p,q\}} = \{11, 10, 01, 00\}$.

Now here is the problem we all face every day. Suppose the agent, a driver waiting to cross the intersection, can see that the light for oncoming cross-traffic has turned red. Semantically, this means the agent only has enough information to rule out states 01 and 00, reducing the possibilities to 11 and 10, one of which must be the actual state of the system. This information can be expressed by the sentence p, since $Mod(p) = \{11, 10\}$.

The agent's difficulty is that the available information isn't really enough for the decision she has to make about driving across — she doesn't know whether the oncoming car will stop, because in state 11 it does and in state 10 it doesn't.

Can reasoning help the agent? Given the available information, p, can the agent figure out whether the car will stop (q)?

Well, it would certainly be nice if it were the case that $p \models q$. The agent knows that p, because she can see it, and if $p \models q$ then she could infer that the oncoming car will stop (i.e. that 11 is the actual state). But of course, things are not that easy: $p \not\models q$ because the state 10, which is a model of p, fails to satisfy q.

On the other hand, if it were the case that $p \models \neg q$, then that would also solve the problem. The agent would infer that the oncoming car is not going to stop (i.e. that 10 is the actual state). However, $p \not\models \neg q$ since 11 is a model of p that fails to satisfy $\neg q$.

So the difficulty facing the agent is that the information available to her in the form of p is insufficient to decide between q and $\neg q$, because neither $p \models q$ nor $p \models \neg q$. This often happens, in other words it is often the case that an agent has too little information for the very strict classical entailment relation $\models$ to be helpful.

What information would the agent need to add to p in order for classical logic to be useful for this bit of everyday decision-making? What information would the agent require in order to conclude that, say, 11 is the actual state?

Well, one possibility is for the agent simply to wait until it is possible to observe whether q becomes true or $\neg q$ becomes true. But that's cheating. In real life, the driver cannot put off acting until she confirms that the oncoming car has stopped. The driver has to decide, with the information available now, whether q or $\neg q$.

Here is another possibility. We know that $p \wedge (p \to q) \models q$, so that if the agent could add, to the observed information p, the rule $p \to q$, then the inference from premiss $p \wedge (p \to q)$ to conclusion q would follow and the driver would know what to do.

But alas. The rule $p \to q$ cannot realistically be adopted as a premiss. The rule does indeed represent a social convention — cars are supposed to stop at red lights — but social conventions can be broken. Unless the driver lives in a dream world, she would in the past have observed someone go over a red light and would therefore know that the sentence $p \to q$ might be false in the actual state of the system. What the driver knows is that cars *normally* stop at red lights, not that cars *always* stop at red lights.

In other words, not only do agents often possess too little information for $\models$ to permit useful inferences to be performed, but the sort of information that the agents would need to add is not an iron-clad universal rule like $p \to q$ but a more complex thing called a *default rule.*

Here are some examples of default rules:

- In summer the weather is typically warm.
- Your phone generally works.
- Birds are normally able to fly.
- If two cities are large and close together, then they are likely to be connected by train.
- In New Zealand, vehicles usually drive on the left.

A default rule is a rule with exceptions. We may consider a default rule as something that expresses *heuristic information*, or a rule of thumb. Drivers know that if the light is red then the system would *typically* be in the state 11, because *typically* the oncoming car stops. But drivers also know that the system doesn't absolutely have to be in state 11 and may instead be in the exceptional state 10, in which case the oncoming car *won't* stop.

Unfortunately there does not seem to be any way to express in $L_{\{p,q\}}$ a rule of thumb saying that something is *normally* the case, e.g. that 11 is a more typical state of affairs than 10.

Since $L_{\{p,q\}}$ cannot express it, we shall represent heuristic information semantically, by making use of preference relations. First we get some technical preliminaries out of the way, and then we'll see how to use such preference relations to formalise everyday commonsense reasoning.

Preference relations

We shall focus on a particular kind of preference relation, called a total preorder.

Definition 1 *The set of ordered pairs $\preccurlyeq$ is a total preorder on S if and only if*

- $\preccurlyeq \subseteq S \times S$, *i.e. if $x \preccurlyeq y$ then $x \in S$ and $y \in S$*
- $\preccurlyeq$ *is reflexive on S, i.e. $x \preccurlyeq x$ for every $x \in S$*

- *$\preccurlyeq$ is transitive, i.e. if $x \preccurlyeq y$ and $y \preccurlyeq z$ then $x \preccurlyeq z$ for all $x, y, z \in S$*
- *$\preccurlyeq$ is total, i.e. $x \preccurlyeq y$ or $y \preccurlyeq x$ or both, for all $x, y \in S$.*

Let's look at an example. The default rule (i.e. heuristic information) that cars normally stop for red lights and normally keep going at green lights may be represented by the total preorder depicted by the diagram, where we take $x \preccurlyeq y$ if and only if x occupies a level below, or at least no higher than, the level of y:

11	00
10	01

Such filing-cabinet diagrams are a convenient way to describe total preorders, because a total preorder $\preccurlyeq$ on a set S always stratifies S into levels. Items x and y will live on the same level if and only if $x \preccurlyeq y$ and $y \preccurlyeq x$.

Our direction of preference goes upward, so the higher an item sits the more it is preferred. Items on the same level are equally preferred. In our traffic example, the preference relation depicts normality: states that are more preferred are more normal. Both 11 and 00 are maximally normal in this example, though they might not be for a different preference relation.

Definition 2 *If $x \preccurlyeq y$ but not $y \preccurlyeq x$ then we may write $x \prec y$. We read $x \prec y$ as "x is strictly below y" or as "y is strictly above x".*

Diagrammatically, if a state x is in a level lower than y, then that tells us that $x \prec y$.

Going back to our example, we see that $10 \prec 11$, so the preorder thinks 11 is more normal than 10. This makes sense because 10 is a state in which the light is red but the oncoming car keeps going instead of stopping, which is certainly a less normal state of affairs than 11, the state in which the light is red and the oncoming car stops. Recall that the waiting driver who sees the light is red (p) knows the actual state is one of 11 and 10 but doesn't know which. The ordering encodes the heuristic information that although this agent (the waiting driver) cannot completely rule out the possibility of 10 being the actual state, the agent knows 11 to be the more normal state of affairs and 10 to be exceptional and unusual.

Definition 3 *Suppose $\preccurlyeq$ is a total preorder on S and let $X \subseteq S$. Call t a maximal element of X if $t \in X$ and there is no $u \in X$ such that $t \prec u$.*

The idea is that if t is a maximal element of X then t is as normal as it gets in X. Recall that in the traffic example, the waiting driver could see that the light was red and therefore knew p. Take $Mod(p) = \{11, 10\}$ to be the subset X. What are its maximal elements? Using the preorder depicted previously, it is clear that 11 is the (only) maximal element in $Mod(p)$. Thus if the light is red (p is true) then the most normal state of affairs is that in which the oncoming car stops (state 11, making q true).

We are going to use such maximal elements to define a new entailment relation. But first, let's make sure we feel comfortable with preference relations.

Remark 4 *There are several things worth noticing about maximal elements:*

- *A subset $X \subseteq S$ may have more than one maximal element, which is why we do not speak of "the maximum" unless we happen to know there is only one.*
- *If t is a maximal element of $X \subseteq S$ then t need not be a maximal element of S as a whole, and more generally t may be maximal in one subset but not in another. We should always mention the subset relative to which an item is maximal.*

- *If t is a maximal element of X, then $u \preccurlyeq t$ for every $u \in X$. (Why? Since $\preccurlyeq$ is total, we have for each $u \in X$ that $u \preccurlyeq t$ or $t \preccurlyeq u$ or both. We cannot ever have that $t \preccurlyeq u$ without also having $u \preccurlyeq t$ because we cannot have that $t \prec u$. So whether or not $t \preccurlyeq u$, we need in every case to have $u \preccurlyeq t$.)*
- *If S is a finite set, then every nonempty subset X of S will have at least one maximal element (the fact that S is finite means X is also finite, and so we can climb up only finitely many levels before we have to stop at the top).*
- *If S is finite and $X \subseteq S$, then for every $x \in X$ we have that either x itself is a maximal element of X or else there is some maximal element t of X such that $x \prec t$. (Think of the maximal elements of X as forming a ceiling, then every element of X either is part of the ceiling or has the ceiling over its head.)*
- *If S is **infinite**, then it is possible for subsets of S to stretch upwards indefinitely without ever hitting a ceiling consisting of maximal elements. In such circumstances, if we wanted to use preference relations as below then we would need to insist on an additional property for $\preccurlyeq$ called smoothness. To avoid this complication, we restrict consideration to **finite** sets of states, so that all nonempty sets of states have maximal elements.*

Where do total preorders come from? Basically any sort of experience that allows the agent to learn something about likelihoods or about what is typical. One source of preference relations is statistical information. Imagine having a team of people stationed at traffic intersections recording information about the frequency with which cars stop (or fail to stop) at red lights. The probabilities generated by these data can be used to order states of the system according to their likelihood. Such a statistical origin is relatively uncommon in everyday life, however, as it requires too much labour to generate the probabilities. Most of our information about typicality is qualitative and gained effortlessly by experience, just as we all know that birds are typically able to fly even though none of us can assign a precise probability to it and all of us are aware of exceptions like penguins or kiwi. You know that your stove usually works, but I'll bet you haven't kept track of the number of times you've switched it on and the number of times it has then worked properly, so your default rule is qualitative rather than statistical. This is why we use preference relations and not probabilities.

A new kind of entailment relation

We are now ready to formalise the kind of everyday commonsense reasoning used by a driver at an intersection (and more generally by agents dealing with everyday decisions).

Assume the agent has a propositional language L_A of the sort we've been using. Assume that the semantics is provided by a finite set S of states and a labelling function $V : S \longrightarrow W_A$ and that $\preccurlyeq$ is a total preorder on S. We now define a new entailment relation for our logic, which we call $\mid\!\sim$ (pronounced "twiddle").

Definition 5 *Let the set of maximal elements in $Mod(\varphi)$ be denoted by $Max(\varphi)$.*
Then $\varphi \mid\!\sim \psi$ if and only if $Max(\varphi) \subseteq Mod(\psi)$.
That is, $\varphi \mid\!\sim \psi$ if and only if every maximal model of φ satisfies ψ.

We see that the main difference between $\varphi \models \psi$ and $\varphi \mid\!\sim \psi$ is that we no longer require **all** models of φ to satisfy ψ, just the maximally preferred models of φ.

Note that if we use a different preference relation we get a different entailment relation $\mid\!\sim$. Since the preference relations are all total preorders, they have lots of properties in common, and below we explore these shared properties. Nevertheless, we should bear in mind that there can be significant differences between the relations $\mid\!\sim$ deriving from differences in the total preorders (e.g. differences in the number of levels into which the total preorder stratifies the set S).

Since our preference relations are total preorders, we call $\mid\!\sim$ a *rational consequence relation*. This name comes from a paper published in 1990 by Kraus, Lehmann, and Magidor. (The full reference is given on the COSC410 coursework webpage — see Notes.)

How does a rational consequence relation $\mid\!\sim$ help us to formalise the commonsense reasoning of the driver at the intersection? Recall that we represented our heuristic information with the aid of $\preccurlyeq$ given by

11	00
10	01

and that the driver could see that the light was red, i.e. could adopt the premiss p. What the driver wondered was whether the oncoming car would stop, i.e. whether q would be true.

Now $p \mid\!\sim q$ since the only maximal model of p is 11 and 11 satisfies q. Hence the driver, if she reasons according to the constraints built into $\mid\!\sim$, can infer that the oncoming car is stopping and will therefore be able to make the decision to drive across the intersection.

It is important to notice that this inference may in fact turn out to be mistaken. It may be that the actual state of the traffic system is 10, because the oncoming car is being driven by someone who has fallen asleep. In this case, the state of affairs is not normal but exceptional, and all bets are off. Sometimes, an agent can do the best she can with the available information and it may still not be good enough. That's real life.

Unlike classical entailment $\models$ in which the movement from premiss to conclusion is so severely constrained that the conclusion is guaranteed to be true in all states making the premiss true, a rational consequence relation like $\mid\!\sim$ requires only that the premiss should make the conclusion reasonable, not that the premiss should guarantee the conclusion. And of course, it is this easing of the constraint that makes $\mid\!\sim$ useful in situations where $\models$ cannot help.

The difference in the thinking governed by $\models$ and the thinking governed by $\mid\!\sim$ is like the difference between mathematics, where every proof establishes a result beyond doubt, and real life reasoning about matters such as where to buy a house, who to marry, what politician to vote for, and whether it's your duty as a parent to insist that your 16-year old son should make his bed every morning. In real life we can't hope to get it right every time — the best we can hope for is to get it right almost all the time.

Nonmonotonic logic

Let us explore the properties of $\mid\!\sim$.

Theorem 6 *$\mid\!\sim$ is reflexive: for all $\alpha \in L_A$, $\alpha \mid\!\sim \alpha$.*

Proof. *Let $\alpha \in L_A$.*

Then every maximal model of α satisfies α, i.e. $Max(\alpha) \subseteq Mod(\alpha)$.

Hence $\alpha \mid\!\sim \alpha$. ■

Not all the properties possessed by $\models$ are necessarily possessed by all instances of $\mathrel{|}\!\sim$. For example, we know that $\models$ is monotonic: if $\alpha \models \beta$ then for any sentence φ, we have $\alpha \wedge \varphi \models \beta$. This need not be the case with $\mathrel{|}\!\sim$.

Theorem 7 *$\mathrel{|}\!\sim$ need not be monotonic.*

Proof. *We shall construct an example of a relation $\mathrel{|}\!\sim$ for which it is possible to find sentences α, β, and φ such that $\alpha \mathrel{|}\!\sim \beta$ but it is not the case that $\alpha \wedge \varphi \mathrel{|}\!\sim \beta$.*

Let $A = \{p, q\}$, $S = \{11, 10, 01, 00\}$ and take $\preccurlyeq$ to be given by the diagram

11	
10	01
00	

As sentences take $\alpha = p$, $\beta = q$, and $\varphi = \neg q$.

Now $p \mathrel{|}\!\sim q$ since $Max(p) = \{11\}$ and 11 satisfies q.

But it is not the case that $p \wedge \neg q \mathrel{|}\!\sim q$ since $Max(p \wedge \neg q) = \{10\}$ and $10 \notin Mod(q)$.

Thus we have found instances of α, β, and φ such that $\alpha \mathrel{|}\!\sim \beta$ but not $\alpha \wedge \varphi \mathrel{|}\!\sim \beta$. ∎

Because of this difference between classical entailment $\models$ and rational consequence relations $\mathrel{|}\!\sim$, the new kind of logic is called *nonmonotonic logic*. The monotonicity of classical logic says that if the agent's store of information grows, then everything the agent previously inferred can still be inferred. Nonmonotonic logic recognises that when we make use of heuristic information to do an inference, we're implicitly assuming everything is normal, so that if we subsequently learn that the circumstances are abnormal we may be forced to retract the earlier inference. Can we somehow distinguish between new information that is safe and new information that tells the agent things are abnormal?

Suppose $\alpha \mathrel{|}\!\sim \beta$ and the agent learns the new information φ. What makes φ potentially unsafe is that it may be undermining α. This needn't be as blatant as contradicting α, and may be as subtle as shifting attention towards the less normal models of α. For example, the driver who was waiting to cross the intersection and observed that the light for cross traffic changed to red ($\alpha = p$) may have inferred that the oncoming car would stop ($\beta = q$) but just as she takes her foot off the brake she hears sirens suggesting that the oncoming car is being pursued by police (φ). This doesn't contradict α, but it does suggest that the situation is not normal, and therefore calls into question the earlier conclusion that the car would stop (β).

If instead of the unexpected new information φ our driver had learnt something that could reasonably have been expected if one knew that α was the case, then the new information should be safe and should not undermine old inferences. This suggests a new property to investigate.

Theorem 8 *$\mathrel{|}\!\sim$ is cautiously monotonic: if $\alpha \mathrel{|}\!\sim \beta$ and $\alpha \mathrel{|}\!\sim \varphi$ then $\alpha \wedge \varphi \mathrel{|}\!\sim \beta$.*

Proof. *Suppose $\alpha \mathrel{|}\!\sim \beta$ and $\alpha \mathrel{|}\!\sim \varphi$.*

Is it the case that $\alpha \wedge \varphi \mathrel{|}\!\sim \beta$?

Pick any $x \in Max(\alpha \wedge \varphi)$. We want to show x satisfies β.

We know that $\alpha \mathrel{|}\!\sim \beta$, so if we can show $x \in Max(\alpha)$ the result would follow.

How can we show that $x \in Max(\alpha)$?

Well, either $x \in Max(\alpha)$ or $x \notin Max(\alpha)$.

Assume that $x \notin Max(\alpha)$ and let's derive a contradiction.

If $x \notin Max(\alpha)$ then there is some $t \in Max(\alpha)$ with $x \prec t$.

But t is a model of $\alpha \wedge \varphi$ since $\alpha \mathrel{|\!\sim} \varphi$

and so x and t are both in $Mod(\alpha \wedge \varphi)$ with $x \prec t$

which contradicts the fact that x is maximal in $Mod(\alpha \wedge \varphi)$.

Hence $x \in Max(\alpha)$ and the desired result follows. ■

Continuing our exploration, recall that classical entailment $\models$ was contrapositive: if $\varphi \models \psi$ then $\neg\psi \models \neg\varphi$. We show that this need not be the case for rational consequence relations $\mathrel{|\!\sim}$.

Theorem 9 *$\mathrel{|\!\sim}$ need not be contrapositive.*

Proof. *We construct an example of $\mathrel{|\!\sim}$ for which there are φ, ψ with $\varphi \mathrel{|\!\sim} \psi$ but not $\neg\psi \mathrel{|\!\sim} \neg\varphi$.*

Let $A = \{p, q\}$, $S = \{11, 10, 01, 00\}$ and take $\preccurlyeq$ to be given by the diagram

11	
10	01
00	

Take $\varphi = p$ and $\psi = q$.

Now $p \mathrel{|\!\sim} q$ since 11 is the only maximal model of p and satisfies q.

But $\neg q$ has 10 as its only maximal model, and 10 fails to satisfy $\neg p$.

Thus it is not the case that $\neg q \mathrel{|\!\sim} \neg p$. ■

Let's look a bit more closely at the family resemblance between $\models$ and the various rational consequence relations $\mathrel{|\!\sim}$. We should realise that every $\mathrel{|\!\sim}$ sanctions all the inferences that $\models$ sanctions, plus usually some more. This makes sense because $\models$ is supposed to be the strictest constraint on moving from premiss to conclusion while $\mathrel{|\!\sim}$ is more relaxed.

Theorem 10 *$\mathrel{|\!\sim}$ is supraclassical: if $\alpha \models \beta$ then $\alpha \mathrel{|\!\sim} \beta$.*

Proof. *Suppose $\alpha \models \beta$.*

Then $Mod(\alpha) \subseteq Mod(\beta)$.

So $Max(\alpha) \subseteq Mod(\beta)$. ■

The supraclassicality of $\mathrel{|\!\sim}$ is what makes it so useful but also where things can go wrong. Inside $\mathrel{|\!\sim}$ sits $\models$, and we know that inferences sanctioned by $\models$ are guaranteed. But around $\models$, yet still inside the supraclassical relation $\mathrel{|\!\sim}$, there sit zero, one, or more additional inferences, and under exceptional circumstances (i.e. in a state that is not maximally normal) any of these additional inferences may turn out to be mistaken. We can say that the inferences sanctioned by $\mathrel{|\!\sim}$ are *defeasible*, because some of those inferences may be defeated by unusual circumstances.

We should also keep in mind that when we generalised from $\models$ to the rational consequence relations $\mathrel{|\!\sim}$, we didn't throw $\models$ away; we merely amplified it. Given the right preference relation $\preccurlyeq$, we can recover $\models$ as one of our rational consequence relations.

Theorem 11 $\models$ *is itself also a rational consequence relation.*

Proof. *All we need to do is pick* $\preccurlyeq$ *so that* $Max(\alpha) = Mod(\alpha)$ *for every* $\alpha \in L_A$.

Let $\preccurlyeq$ *be the total preorder* $S \times S$.

Now all states are on the same level, i.e. no state is strictly below another.

Hence every model of α *is maximal in* $Mod(\alpha)$.

Thus $Max(\alpha) \subseteq Mod(\beta)$ *if and only if* $Mod(\alpha) \subseteq Mod(\beta)$. ∎

The rest is left to the exercises.

Exercises

Quiz: The quiz question for lecture 5 will come from the last section, dealing with important properties of rational consequence relations, and will involve one of exercises 1, 2, or 3 (i.e. the properties of **explosiveness**, **and**, and **or**).

Using total preorders to represent default rules

1. Suppose we are modelling a new kind of helicopter, which has both a rotor and a jet engine. As representation language we use $L_{\{p,q\}}$ where p stands for "The jet is on" and q for "The rotor is on". The possible states of the system are $S = W_A = \{11, 10, 01, 00\}$.

   Give a total preorder (by means of a filing-cabinet diagram) that represents the following heuristic information.

   - It is very unusual for the jet to be on while the rotor is off.
   - It is normally the case that the helicopter is on the ground with everything switched off.
   - It is less normal but not very unusual for the helicopter to be flying, in which case it is equally likely to have just the rotor on as it is to have both rotor and jet on.

2. Suppose discovery of a cheap non-polluting source of fuel means that the helicopter can spend most of its time in the air, and that the rotor is needed only for taking off and landing. (Thus the jet alone is used in the air, while the jet and rotor together are used for take-offs and landings. The rotor is never used alone.)

   Assume that the helicopter spends about as much time on the ground undergoing maintenance as it does doing take-offs and landings. (But remember that most of the time is spent flying, i.e. not on the ground and not in the process of taking off or landing.)

   Construct a new total preorder depicting this heuristic information.

3. Consider the 3 Card System in which each of three players get one of three cards coloured red, green or blue. Suppose the representation language has 9 atoms r_1, r_2, r_3, g_1, g_2, g_3, b_1, b_2, b_3 where r_1 stands for "Player 1 has the red card" and so on.

   The system has six possible states: $S = \{rgb, rbg, grb, gbr, brg, bgr\}$ where rgb is the state/deal giving the red card to player 1 (so that r_1 is true, r_2 and r_3 false), the green card to player 2 (so that g_2 is true while g_1 and g_3 are false), the blue card to player 3 (so that b_3 is true while b_1 and b_2 are false) and so forth.

Now imagine that the cards are dealt to players by a new method.

- The dealer rolls a 6-sided die, each side of which is equally likely to be uppermost.
- If the uppermost face shows a 1, 2, 3, or 4, the dealer gives the red card to player 1.
- If the uppermost face is 5, player 1 gets the green card.
- If the uppermost face is 6, player 1 gets the blue card.
- Next the dealer shuffles the remaining two cards and gives one to player 2 and the last to player 3.

Construct a total preorder depicting the relative likelihoods of the six states, with more probable being more preferred. You do not need to calculate precise probabilities for this.

Examples of rational consequence relations

1. Recall that we previously modelled a new kind of helicopter by constructing a total preorder that portrays the following rules of thumb. It is very unusual for the jet to be on while the rotor is off. It is normally the case that the helicopter is on the ground with everything switched off. It is less normal but not very unusual for the helicopter to be flying, in which case it is equally likely to have just the rotor on as it is to have both rotor and jet on.

   Let $\mid\!\sim$ be the rational consequence relation induced by this preorder. Recall that p expresses that the jet is on and q that the rotor is on. Does

   - $p \mid\!\sim p \wedge q$?
   - $q \mid\!\sim p \wedge q$?
   - $p \vee q \mid\!\sim p$?
   - $p \vee q \mid\!\sim q$?
   - $p \vee \neg p \mid\!\sim p$?
   - $p \vee \neg p \mid\!\sim \neg p$?

2. Recall that discovery of a cheap non-polluting source of fuel changes things so that that the helicopter can spend most of its time in the air, the rotor is needed only for taking off and landing, and the helicopter spends about as much time on the ground undergoing maintenance as it does doing take-offs or landings. If $\mid\!\sim$ is the rational consequence relation induced by the new total preorder, is it the case that

   - $p \mid\!\sim p \wedge q$?
   - $q \mid\!\sim p \wedge q$?
   - $p \vee q \mid\!\sim p$?
   - $p \vee q \mid\!\sim q$?
   - $p \vee \neg p \mid\!\sim p$?
   - $p \vee \neg p \mid\!\sim \neg p$?

3. Recall the 3 Card System in which three players each get one of three cards coloured red, green or blue. And recall that the cards are dealt to players by a new method in which the dealer first rolls a 6-sided die each side of which is equally likely to be uppermost. If the uppermost face shows a 1, 2, 3, or 4, the dealer gives the red card to player 1. If the uppermost face is 5, player 1 gets the green card, and if it is 6 he gets the blue card. The dealer then shuffles the remaining cards, giving one to player 2 and the last to player 3.

   Let $\mid\sim$ be the rational consequence relation induced by the total preorder that regards the most probable state as the most normal and less probable states as less normal. Is it the case that

   - $g_2 \mid\sim r_1$?
   - $g_3 \mid\sim r_2 \vee b_1$?
   - $r_1 \mid\sim g_2 \rightarrow b_3$?
   - $r_1 \mid\sim g_2 \vee b_2$?

Important properties of rational consequence relations

Let L_A be arbitrary, and let $\mid\sim$ be an arbitrary rational consequence relation for L_A. For each of the following properties, decide whether $\mid\sim$ has the property. If so, prove it. If not, give a counterexample. (Exercises 10, 11, and 12 are quite challenging but do-able with what you know.)

1. (Explosiveness) $\alpha \wedge \neg\alpha \mid\sim \beta$ for every $\alpha, \beta \in L_A$.
2. (And) If $\alpha \mid\sim \beta$ and $\alpha \mid\sim \varphi$ then $\alpha \mid\sim \beta \wedge \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
3. (Or) If $\alpha \mid\sim \beta$ and $\varphi \mid\sim \beta$ then $\alpha \vee \varphi \mid\sim \beta$ for all $\alpha, \beta, \varphi \in L_A$.
4. (Right weakening) If $\alpha \mid\sim \beta$ and $\beta \models \varphi$ then $\alpha \mid\sim \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
5. (Transitivity) If $\alpha \mid\sim \beta$ and $\beta \mid\sim \varphi$ then $\alpha \mid\sim \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
6. (Cut) If $\alpha \wedge \beta \mid\sim \varphi$ and $\alpha \mid\sim \beta$ then $\alpha \mid\sim \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
7. (Conditional insertion) If $\alpha \wedge \beta \mid\sim \varphi$ then $\alpha \mid\sim \beta \rightarrow \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
8. (Conditional elimination) If $\alpha \mid\sim \beta \rightarrow \varphi$ then $\alpha \wedge \beta \mid\sim \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
9. (Classical equivalence) If $\alpha \equiv \beta$ and $\alpha \mid\sim \varphi$ then $\beta \mid\sim \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
10. (Rational equivalence) If $\alpha \mid\sim \beta$ and $\beta \mid\sim \alpha$, and $\alpha \mid\sim \varphi$, then $\beta \mid\sim \varphi$ for all $\alpha, \beta, \varphi \in L_A$.
11. (Rational monotonicity) If $\alpha \mid\sim \beta$ and it is not the case that $\alpha \mid\sim \neg\varphi$, then $\alpha \wedge \varphi \mid\sim \beta$.
12. (Rational or) If $\alpha \mid\sim \beta$ then either $\alpha \wedge \varphi \mid\sim \beta$ or $\alpha \wedge \neg\varphi \mid\sim \beta$.