

Chennai Worlds 2014 Adjudicator Briefing

Contents

- I. Deciding Who Wins
- II. Decision-Making Process
- III. Deliberation
- IV. Announcing the Result
- V. Giving Feedback
- VI. Some Pitfalls to Avoid in Decision-Making and Feedback

I. Deciding Who Wins

Objective of teams

Judges assess which team did the best job at **persuading them** that the motion, in the context in which it was discussed during the round, ought to be affirmed or rejected. Judges do this as an **'informed global citizen'** and consider the strength of the arguments that each team presented.

Criteria

How one determines what constitutes a good/bad policy is itself debatable. Teams are allowed to debate the criteria for assessing whether a policy is good as part of arguing whether it is, in fact, good. Judges should adjudicate this debate about criteria, as opposed to just applying their own. Judges should thus listen to teams' arguments about what our aims and principles should be, and evaluate the claims of harms or benefits in that context.

This can make these claims particularly vital; if in the debate 'THW Invade North Korea' you show, for instance, 'war is always wrong, regardless of the practical benefits', the other side's claims about the practical benefits of invading North Korea are irrelevant until they've proved that a practical calculus is relevant.

Note that there is no sense in which this argument is too 'generic' to gain credit. There is an absolute requirement that arguments be relevant to a debate. There is absolutely no requirement that arguments pertain only to that debate.

For instance, in a debate about the motion 'THW ban all forms of prostitution,' an argument from the opposition that a black market would be created by enacting this policy may also be an argument that could be used in a debate about banning gambling or alcohol. This is not a flaw in the argument, and its contribution to the debate should be judged by taking into account the quality of the explanation and how other teams dealt with the argument.

Assessing Argumentation

The outcome of the debate should depend on what the teams say. Judges must not intervene in the debate. Do not invent arguments for teams, do not complete arguments, and do not rebut arguments.

We do not consider claims invalid just because we disagree, or because we can see weaknesses in their arguments. Arguments stand as 'true' once they are made and substantiated until they are responded to.

This has an important implication: if OG, for instance, make arguments where the conclusion is 'we should do the policy,' and every other team ignores those arguments, then OG do not lose because 'the debate moved on from them'. Rather, their unrebutted arguments should still be treated as true and should be weighed as such. That does NOT mean that the unresponded-to arguments have a particular effect on the ranking of OG in this example. Judges still need to consider how important an argument is before deciding how it affects the ranking of teams in a debate. Indeed, there are some rare cases where a statement or argument being "true" could even harm the team that argued for it.

Responsiveness and The Substantiation of Arguments

If an argument is clearly absurd (such that you cannot conceive of anyone we would consider sane believing its logic and/or premises), or it was of marginal importance to the speech of the speaker making the claim, then it is reasonable for a responding team to decide to spend their time elsewhere, particularly where there is other stronger material in the round. Furthermore, judges are entitled to assess how well substantiated an argument is – an argument that is a just an assertion ("as we'd all agree, language constructs reality") without any subsequent substantiation should not receive much credit, though it may be advantageous for other teams to point out and respond to weakly constructed arguments.

In an ideal world, teams will engage in extensive responses to each others' well-detailed points. In many debates, however, teams will often talk past each other and leave each others' points unchallenged. Under those circumstances, the judge will have to assess not only which arguments are most important, but also which ones are most clearly proven. Points that require the judge to make numerous logical leaps are better than no points at all but are not preferable to a well-reasoned argument that rests on fewer unsubstantiated assumptions.

Understanding which arguments are and are not rebutted is important to judging debates. Note that speakers don't have to use the word "rebuttal" to respond to an argument. It may be tidier if they do, but judges should not ignore material that adequately deals with an argument just because the speaker does not point out that it does. If judges did not follow this rule, it would be very difficult to compare opening and closing half teams. An opening half team may have presented arguments that adequately deal with the closing half's contribution, but it is highly unlikely that they will have labeled it as rebuttal ahead of time.

Speakers should also not be 'punished' for not refuting everything: some claims do not do any harm at all to the opposite side. For example, in a debate about the legalisation of drugs, if the government say 'pink elephants are cute because they have those nice ears and are a pleasant colour', this flawed argument can be safely allowed to stand as it isn't a reason to legalise drugs. There is, therefore, no need to point out that blue elephants are obviously more tasteful.

So too, if they said 'some drugs are less harmful than others', this may also be ignored. While it is clearly more related to the debate than the cute pink elephants argument, it is pre-argumentative – that is, it has not got sufficient surrounding words to be a reason to do

or not do the policy. Opp can quite happily say 'yes, some drugs are more harmful than others' and move on. Or just ignore this argumentative non-sequitur. Often as a judge, it can be tempting to complete arguments that are interesting but pre-argumentative. Don't.

The Importance of Arguments

What matters, once an argument is made, is how important its conclusion seems to be in the debate, based on what people have said in this debate, and the extent to which it seems to be successfully responded to.

We should not consider how important we think a particular argument might be, in the abstract, but rather how important it became in this particular debate. Similarly, we should not care how interesting, or sophisticated an argument was, except in so far as those things tacitly influence us to find those claims to be more true or important and have impacts on how well others can and do respond.

Things that do not matter

Certain things do not matter (in themselves) in debating:

- The number of arguments a team makes,
- How clever or sophisticated the argument was,
- How interesting the argument was.

None of the factors above should affect where a team places in a debate.

To sum up thus far: the team whose arguments are most effective at persuading the judge that their side of the motion is correct wins the debate. This must be assessed solely through material provided in the debate. Don't penalise teams for failing to make arguments. Don't consider arguments that you're aware of but which weren't made. For instance: in the motion "This house would place sanctions on Israel", you may feel that discussing Israel's nuclear policy while not even mentioning the Israeli-Palestinian conflict was a bad strategy on behalf of proposition and that proposition missed the chance to make a lot of good arguments. That may be so, but you may not allow this to affect your ranking.

Mechanism for persuasion

Many debating coaches decompose persuasiveness into different categories, which they then train teams to get better at. It is important to recognise that these are different mechanisms for persuasion, but they are means to the same end.

Examples of such mechanisms are:

- Good analysis
- Good rebuttal
- Effective manner/good style
- Good strategic decisions
- Using examples

If these are done well, they might influence you to believe a team is making the most substantial contribution to the debate. That's great for the team in question. But note that you shouldn't also consider them as separate categories for assessment. Doing so would be double counting them. While it may be helpful in coaching to dissect debating and

examine these factors separately, as judges we should just see what persuasive logic seems to have been produced.

One thing that is certainly unpersuasive is inconsistency. Teams should not contradict themselves or their bench partners (on either government or opposition). Internally inconsistent teams cannot simultaneously get credit for two areas of mutually exclusive argument. We're forced to credit them for the least effective of the two.

Beyond being unpersuasive, inconsistency is unfair to opposing teams. It cannot be reasonably expected from a debater to answer two contradicting lines of argumentation, especially if those are given in different times during the debate.

Judging as an informed global citizen

There is obviously a need for judges to abstract themselves from their personal positions when in debates. If we judged as ourselves then the result would depend strongly on what we know, how we think, and what opinions we hold. We would be making personal judgments about how compelling we find particular arguments.

Therefore, we attempt to judge as an abstract 'informed global citizen' and minimise the impact that our knowledge and opinions have on the result. This is a difficult process, and something that bears thinking about.

It is important, however, to remember that thinking 'as' an 'informed global citizen' does not absolve us from our responsibilities to actually judge the debate – that is to evaluate the logical flow of arguments and determine the extent to which teams have seemed to win them.

An informed global citizen:

- Attempts to evaluate logically what the best thing to do is,
- Uses as inputs the arguments made by the teams,
- Has the sort of knowledge of someone who reads the front pages and world section of a high quality newspaper regularly,
- Has a grounding in world history,
- Is NOT a specialist in any subject, and thus will not be familiar with technical vocabulary or specialist literature,
- Comes from nowhere in particular. There are no 'domestic examples' requiring less explanation, even where everyone in the room comes from that country. Wherever you are from, assume your audience is from somewhere else.

The informed global citizen is, in short, a smart person who has a good deal of knowledge that is broad rather than deep. Imagine a bright and well-read university student who is studying a subject completely alien to the one that is being debated.

Giving credit where it is due

We'd like to say something about awarding credit to teams for the material they bring to the debate. In particular, we want to answer the question "what counts as new material?"

Anything teams say that is a new contribution to the debate is worthy of credit. This means they get credit for everything beneficial they do: arguments, rebuttals, framing, examples, etc. Their contribution is valuable insofar as what they say benefits their side more than

what has already been said. So, a whole new argument that is entirely theirs is valuable to the extent that argument helps their side; an improvement on an existing argument is valuable insofar as their addition to the argument helps their side; rebuttal is valuable insofar as defeating that specific material helps their side. All this is determined by the current state-of-play in the round.

One conclusion of this is that it is not enough for a closing team to make an argument somewhat better than their opening: they gain credit only for the improvement on the previous argument. They must make improvements and additions to the opening case that outweigh the totality of that case in order to place ahead of the top half team on their side of the bench.

Comparing contributions

Where teams have a chance to rebut each other, assessing relative contribution is easy. Judges should track the argument and assess, given their responses to each other, which team's contribution was more significant in furthering their cause to logically persuade us that we should do the policy, or that we should not.

Where teams don't get a chance to rebut others, this is trickier. This happens fairly often:

- Teams on diagonals
- Opposition whip explains something in a new way
- Opening teams shut out of POIs

In these circumstances judges are forced to perform some more independent assessment of the strength of the arguments teams made. The question to ask is 'how strong to potential rebuttals was a team's contribution?'

Ideally, assessing strength will involve a comparison with material on the table, or very minor extensions thereof. For instance, when comparing a closing half team with a team that was diagonally across from them in the front half, judges should first ask whether there was anything in the top half team's case that is relevant or (ideally) rebuttive to the bottom half team's argument. Did the back half team deal with this material? Check whether each team allowed the diagonal team in on POIs to give them an opportunity to engage. Deliberately shutting out engagement from a team whose material is relevant is often obvious and unpersuasive.

Judging Points of Information

Speakers are expected to accept at least one point of information per speech. Points of Information are important in comparing teams and a component in assessing engagement.

A speaker that fails to take a POI (or a point of clarification) during their speech should be punished by judges if they were offered several in the last minute of unprotected time (between 5.00-6.00mins). Failure to take a POI generally indicates a reduced level of engagement. While this does NOT mean that a team will take an automatic fourth for failing to take a POI or that a team that didn't take a POI cannot possibly win a debate, it can (and should) be an important factor in close calls.

Thus, in a very close debate, if the judge is deciding between two comparably matched teams, we would generally expect a team that failed to take any POIs to place behind a team where each speaker accepted at least one POI. If a team has clearly won a debate,

however, the failure to take a POI should NOT change the team rankings. It may, however, reduce the margin of victory.

Please note that while there is NO requirement to allow the person offering a POI to speak for 15 seconds, cutting off the offeror of a POI before the point could reasonably be understood should be treated by the judges as the equivalent of not having taken the POI.

II. Decision Making Process

Once the debate has finished, the judges should decide the rankings of the teams according to the following rules:

- Judges assess which team had the most persuasive speeches in favor of showing that, on balance, their side of the debate is correct. The judges do so as an 'informed global citizen,' applying a very small set of technical rules.
- There is a 20-minute time limit on deliberations.

Decide the ranking first, with no consideration of speaker marks until this has been established. This reflects the fact that teams win debates, not speakers, and they win based on their aggregate contribution. We are not evaluating our aesthetic appreciation of the speeches (or proxy-marking 'team balance'): we're assessing the team's aggregate contribution. Imbalance within a team should be reflected by giving the speakers different speaker marks.

If at the end of the time limit on deliberations a vote is needed, that vote should aim to establish a majority in favor of a call, with the chair holding a casting vote in case of an even split. In the case that there is absolutely no agreement, all judges should vote rank-by-rank from first to fourth, and the teams should be ranked according to the average this method produces. Again, in the event of a tie, the chair holds the tie-breaking vote.

Note, the team ranking should reflect cumulative speaker scores of each team. e.g. no low-point wins.

III. Deliberation

In close rounds, we expect that the judges on the panel will have different views on the debate. Here we sketch some suggestions for how this could be managed.

After a reading of notes, the chair should begin by asking each wing to give a ranking of the teams. This is not binding; it is a working hypothesis that will evolve as the discussion progresses. Wing judges should not feel any pressure to agree with one another or the chair in their initial call, as there is no negative implication for changing your call.

The chair should then assess the level of consensus which exists. There are thousands of possible combinations, but thankfully a few scenarios crop up fairly often.

- Everyone has exactly the same rankings – have a brief discussion to ensure rankings are the same for similar reasons. Move on to scoring.
- Everyone has the same except 1 person – ask him to defend his position. Be specific, tailoring the requested defence to the difference between the minority and majority opinion. If it is a difference of one team, focus on that team, etc.
- There is similarity in rankings (you agree on where 1 team is ranked or some relative rankings – everyone agrees OG is better than CG) but also some crucial differences. Begin by establishing which discussions need to happen (i.e. there is disagreement about whether OO beat OG). Begin by consolidating the consensus that exists, and use this as a platform to break deadlocks.
- Chaos. There is no similarity between the rankings. Guide a discussion of each team's arguments, or, depending on what makes sense to you and in context, of the clashes between particular pairs of teams. These debates often hinge on how one argument was evaluated, so your aim is to detect such differences in interpretation. The initial discussion is intended to inform each other of your perspectives and find some level of common understanding. If two judges believe different arguments are central, frame a discussion about their relative priority. Get each judge to explain their position, and attempt to establish a metric for the importance of arguments in the debate. After this brief discussion, rank the teams and compare again. If you have achieved some overlap, move on to the suggestions under (c) above. Vote if necessary.

Once a ranking has been decided upon, the chair should lead a quick determination of speaker scores. These should reflect the majority decision and should not be a compromise between various opinions. i.e. don't say "We all think OG wins, but we can make sure the speaks reflect your view". If the majority doesn't think a relative ranking is close, there is no reason that the speaker marks should make the round appear closer than the majority thought it was

Achieving a consensus is ideal, but it is not always possible. Opinions may not change or the time it would take to change them is longer than the 20 mins allocated. A split may at some points be a more accurate evaluation of what happened in the debate. Do not make decisions based on shabby compromises, but do not fear to call a vote on issues. During feedback, we expect you to explain the decision to use votes to the debaters and how the out come of these votes affected the final call.

IV. Announcing Results

The chair of the panel delivers the adjudication speech. In the case that the chair loses a vote, she may, if she chooses, retire from this position and ask one of the wing judges who voted in the majority to deliver all or part of the adjudication. If the chair does give the adjudication, this must be to defend the majority position, although the chair should overtly state she disagreed with the majority.

The adjudication speech should distinguish between the reasons for the decision and advice for teams. The reasons should be about what did happen, while advice is about what didn't happen. The latter cannot be a basis for the former.

The primary aim of an adjudication speech is to convey to the teams the reasoning of the panel in ranking the teams as they did. The speech should therefore present a logical argument for the ranking, using as evidence the arguments made in the debate and how they influenced the judges. Debates shouldn't be judged according to coaching models (either prescriptive models like 'problem/solution' or decompositions of persuasiveness like 'content, style, strategy').

The speech should be structured as followed:

- Announce the ranking of the teams.
- Go through the teams in an order that makes sense, comparing pairs of teams and explaining why one beat the other. Typically, this will mean going chronologically (beginning with OG and ending with CO) or in the order of the decision (starting with either the first or fourth-place team).
- Summarise and invite teams for more detailed feedback.

Comparing teams involves more than making isolated statements about Team X and Team Y, and saying 'so X clearly beat Y'. It requires that you explain the interaction between the teams to establish who had the better argument.

Be specific and be detailed – the vague application of adjectives is not sufficient judging. Identify arguments, whether and how they were responded to, and what the impact of the remainder was. Identify which teams get credit for what, and how this influenced your decision about whether or not we should support the motion. One effective way to give feedback on an argument or area of clash is to discuss the contribution of each team on that point in chronological order. In other words, discuss first the contribution made by first proposition, then the contribution made by first opposition, and then explain why one was more persuasive than the other and the factors that went into that decision. Judges are not required to follow this format, but they are required to be comparative and specific.

Advice should be separated from the reasons for your decision, lest you confuse teams about which is which. There are a number of broad areas of advice you may want to give as a judge:

- General advice on how to improve
- Suggestions of reasons why things identified in the adjudication happened
- What might have been run (although minimise this unless asked!)

VI. Some Pitfalls to Avoid in Decision-Making and Feedback

What follows is a common set of mistakes that judges may make in determining results and giving feedback.

Dealing in generalities rather than specifics

- “We thought that second opposition really brought the case home for us, so they won the debate.”
- “First proposition talked about rights, but I really didn’t find it persuasive.”
- “First opposition had some interesting things to say, but the analysis didn’t get better until second opposition.”

It’s perfectly fine for adjudicators to use general language to introduce their reasons, provided that each general statement is supported by examples of what actually happened. No statement of the sorts that we’ve listed above should ever go unsupported by specific examples of the claim being made, either during the deliberation or during feedback.

Failing to judge the debate as it happened

- “Proposition never talked about rights in this debate.”
- “It took until the summation speaker until we heard anything about rights.”
- “I really wouldn’t have propped it like that.”

Judges may have their own opinion as to what the best arguments for each side in the debate will be, but these are not the criteria on which the debate is to be judged. Judges may advise teams that there were interesting avenues of analysis left unexplored, but they may not penalise teams for their approach to the motion, or the things that each team decided to emphasize.

Granting certain “classes” of arguments undue priority

- “Second opposition were the only team who told us about the importance of principles.”
- “Only first proposition knew the names of major Brazilian cities.”
- “Second prop won because their arguments were moral rather than practical.”

This judging pitfall takes a number of forms. One of which is the fetishisation of the use of specific knowledge in the making of arguments. Teams which make strong arguments buttressed by good knowledge should be rewarded, but not because of the total amount of facts they named, but because of the strength of the arguments which those facts were marshaled in support of. A clever use of facts makes an argument stronger and better; it does not make an argument.

A second form of this pitfall is according improper priority to arguments that are of various types (moral/philosophical/economic/practical). A ‘principled’ argument, for example, is not necessarily better or worse than a practical one. It depends what each argument seeks to prove and how well it does so.

'Penalty judging'

- "Your last point came after six minutes in your speech, so that really hurt your team."
- "We had questions about the mechanism, so we put you last."

A good judge isn't one who tries to find as many reasons as possible to exclude consideration of a team's arguments and speak instead about the form- rather than the content- of their contribution. If a team violates the duties of role fulfillment, they should be penalised only up to the point of removing any harm they caused to the debate through failure to fulfill their role. The one exception is a failure to take a point of information, which is discussed above.

Some Examples:

- Beginning a point after six minutes probably means a speaker will have less time to develop it, but a judge should still evaluate how substantial the argument's contribution is to the round. A one-minute argument can be just as persuasive in the last minute of a speech as it can somewhere in the middle.
- Not taking any POIs means that a speaker's material is to be viewed as less persuasive, not excluded from consideration.
- Lack of clarity in a mechanism should be resolved by allowing the opposition teams to make any reasonable assumptions of their own and letting the debate carry on from there. It may also make Opening Government's case less persuasive if the lack of clarity in the mechanism makes it seem less plausible that the policy could be carried out or if the ambiguity calls into question some of the policy's benefits.
- If a speaker introduces new arguments in an opposition summation speech, these are to be discounted, as though the speaker had said absolutely nothing.

To put it a different way, a lot of mistakes that judges "penalise" in debates are really just instances where the speaker's decisions have created an opportunity cost. The speaker could have used her time more effectively, perhaps by establishing a clearer mechanism, but her failure to do so will make her speech less persuasive already; there is no need to add an additional "penalty" by double counting the mistake.

Judging on Format Rather than Content

- "You only spoke for one minute about the effect of this policy on Wales"
- "You should have put your argument about rights first."
- "Your team was unbalanced- all the good points came from the first speaker."
- "You only spoke for five minutes."

Speaking for a certain length of time or placing arguments in a certain order is irrelevant (in and of itself) to which team won the debate. Naturally, speakers and teams who spend all their time on good arguments and spend more time explaining more important and more complex arguments will do better at being persuasive, but they succeed because they have made good arguments and have explained those arguments well, not because they 'spent time on them.' A speaker can win a debate with a one minute speech (but it's very, very hard to do so). Similarly, sometimes it will make a speech more persuasive to discuss arguments in a particular order because a later claim builds on the analysis of an earlier claim. But if two arguments are genuinely independent, it should not matter in what order a speaker chose to deliver them.

With thanks to the DLSU 2012 and Berlin 2013 Adjudication teams, who drafted the original versions of this guidance. Our updates have been minor in comparison to their extensive work.