

The Only Thing We Have to Fear:

Tribute's first 6 months

Brandon Fogel, September 2020

In “To Whom Tribute Is Due” (Diplomacy World #149, April 2020), I tried to recast Diplomacy scoring system analysis in terms of general incentives rather than specific outcomes. In the last section of the paper, I described a new scoring system borne out of that analysis, Tribute, which was developed collaboratively over 2019 by a cadre of Windy City Weasels, namely Chris Kelly, Bryan Pravel, Jake Trotta, and myself. The Weasels began implementing Tribute in league play in February of 2020, and since the April article, several other competitions have adopted it.

Tribute has produced a wide variety of reactions, some positive, some not. Diplomacy players tend to get attached to certain styles of play, often those promoted by the scoring system they learn first, and they can go further to develop views about what counts as “natural” or “pure” Diplomacy. This can produce strong reactions to scoring innovations, and Tribute has not escaped this fate. Here I will review the system's performance to date, respond to some criticisms, and propose a possible modification.

1. The big picture

1.1. Tribute

In Tribute scoring, every game is worth 100 points. If a player solos, they get all 100 points. In the event of a draw, each surviving player gets 1 point for each supply center plus an equal share of the remaining 66 points, and for each center the board-topper has over 6, the non-toppers pay 1 point in tribute to the board-topper. Shared board-toppers split the tribute paid by the non-toppers equally. 6 was chosen as the tribute threshold because it's the smallest possible lone board-top.

Players sometimes try to understand Tribute as a variation on other scoring systems. “It's Draw-Size with a topping bonus.” “It's C-Diplo with a survival bonus.” Tribute does have commonalities with both of those systems, but it also represents a radical departure. Because the survival bonus is in competition with the topping bonus, every surviving player has a stake in what happens with the board-topper's center count, regardless of their size. Likewise the board-topper has incentive to grow as big as possible, not just to top the board. Neither is the case in Draw-Size or C-Diplo.

1.2.A note on strategy under Tribute

Tribute is primarily a board-topping system. There are rewards for survival and to a lesser extent merely for accumulating centers, but the way to score big in Tribute, short of a solo, is to top the board. Players who play to reduce the draw size without aiming to top the board will likely end up with mediocre scores, not terrible but also not good.

Because Tribute incentivizes survival, it also necessarily incentivizes eliminating other players. However, the reward for eliminating another player is relatively small, just 1-3 points for the first few eliminations. The total increase in survival bonus for getting from 7 players to 4 is 7.1 points. Not insignificant, but dwarfed by the benefit of topping the board, which is usually in the range of 20-40 points. Furthermore, focusing on eliminating other players may simply allow the board-topper to grow bigger, reducing the potential reward.

As I've written previously, players should aim to top the board, up to the point that there really is little hope of achieving it. In that case, players should aim to survive and keep the board-topper as small as possible.

1.3.Incentive measures

In the original paper, I provided charts that measure the strength of the chosen incentives as functions of other variables, like center count or rank. That was useful for seeing some of the details about system behavior relative to context. Now I'd like to consolidate those measures into single numbers. Note that I have added the Dixie scoring system to these calculations. Dixie is a variant of Draw-Size scoring that awards some points to eliminated and surrendering players based on order of elimination.

Growth incentive

System	Value
Tribute	4.65
Sum of Squares	4.35
C-Diplo	4.01
Carnage	1.44
Dixie	1.03
Draw-Size	0.53

Average change in score for taking 1 dot, averaged first by center count.

Survival incentive

System	Value
DSS	16.65
Dixie	11.90
Tribute	5.82
C-Diplo	1.01
Sum of Squares	0.26
Carnage	0.00

Average value of surviving as the smallest power with a single dot.

Board-Top incentive

System	Value
Tribute	21.55
C-Diplo	20.49
Sum of Squares	5.97
Carnage	3.53
Dixie	0.68
Draw-Size	0.00

Average change in score for taking the lead, averaged first by center count.

Dominance incentive

System	Value
Sum of Squares	5.65
Tribute	5.42
C-Diplo	1.00
Dixie	0.95
Draw-Size	0.39
Carnage	0.00

Average change in score for growing by 1 dot as the board-topper, averaged first by center count.

Balance of Power incentive

System	Value
Tribute	27%
Sum of Squares	11%
C-Diplo	-2%
Dixie	-9%
Draw-Size	-12%
Carnage	-15%

Degree to which taking a dot from the leader is better than taking one from another player, as a percentage of the theoretical maximum, averaged first by rank.

Tribute scores well on each incentive measure, which is no accident. We chose these incentives because we value them, and then we developed a scoring system specifically to maximize them. Whether these incentives are good ones to focus on is up for debate, as is whether these are good measures of those incentives.

These measures do have a shortcoming that is worth some discussion. They are all calculations of “next dot” changes in score, which means they only measure the short term impact of a particular dot exchange. Of course, long-term strategic thinking is important in Diplomacy, and ideally the incentive measures would take that into account.

As of yet I haven't found a way to tame the computational complexity of multiple-dot exchanges.

The measure of balance of power incentive probably suffers most in this regard. I can think of three reasons a player might be motivated to fight the leader rather than another player: a) because board-topping is valuable, b) because the leader growing will hurt them, or c) because taking a dot from the leader is worth a bit more than taking one from someone else. I suspect all are important, meaning a good measure should account for all three, whereas the current measure only looks at c.

2. Results

Along with the current Weasels league season, Tribute has been adopted by the Virtual Diplomacy League (an online face-to-face competition), Nexus Season 5 (an online full press extended-deadline competition), Nexus's inaugural gunboat tournament, and Weasel Moot XIV (Chicago's annual face-to-face tournament, online this year), as well as a series of private gunboat tournaments that I run. Here are some statistics from those competitions along with some others for comparison.

In the tables below, the "Survival Value" of a game is calculated by taking the smallest score among the surviving players and subtracting the largest score among the eliminated players. Also, scores for all systems have been normalized so that the total score for each game is 100.

2.1. Time-unlimited face-to-face

We have had three tournaments with time-unlimited games so far this year, all virtual. Here are the compiled statistics from those tournaments:

Competition	System	Games	Avg Years	Avg Draw Size	Shared Tops	Avg Topper Score	Avg Survival Value
DixieCon	Dixie	20	7.9	3.8	20%	26.2	17.8
Massacre	Carnage	14	7.9	5.6	21%	24.6	4.1
Weasel Moot	Tribute	14	9.9	4.1	57%	39.3	19.4

Here Tribute shows marked differences from the other scoring systems. Games went significantly longer—2 game years on average—and resulted in dramatically more shared tops, especially in the first round at Weasel Moot, where 5 of 7 games had shared tops (including a 4-way). Less surprising is that Tribute resulted in more eliminations than in games under Carnage, where there is no survival or elimination incentive. Perhaps a little

surprising is that survival was worth a bit more under Tribute than Dixie, which is often mistaken for simple Draw-Size, but this is may just be due to small sample size.

The high number of shared tops is both surprising and fascinating. Discussions with players revealed that the presence of the top board at Weasel Moot had a strong effect on play, especially in the first round, where 5 of 7 games ended in shared tops. Many players said they were content with scores in the 20s, knowing they only needed to make the top 7 and also wanting to avoid becoming a target in the second round. This logic would likely have been attractive under other scoring systems as well, although Tribute certainly provides bigger temptation than other scoring systems to subvert an agreement to arrange a shared top. Whether sticking to such agreements is a winning strategy under Tribute in the long run is up for debate.

2.2. Time-limited face-to-face

Time-limited games play differently from time-unlimited games, since players don't need to bring the game to a point where everyone agrees the game isn't worth continuing. One would expect fewer eliminations and perhaps lower scores for board toppers in systems like Tribute and Sum of Squares, since there is less time to amass a higher center count.

Competition	System	Games	Avg Years	Avg Draw Size	Shared Tops	Avg Topper Score	Avg Survival Value
Weasels	Tribute	13	6.5	5.5	15%	41.9	10.1
VDL	Tribute	13	7.7	5.3	31%	38.6	11.1
VDL	SoS	7	6.9	5.4	0%	47.7	3.9
Liberty Cup	Carnage	19	6.8	4.9	26%	24.4	4.2

VDL has included both Sum of Squares and Tribute games this season, providing a particularly interesting comparison. Tribute has resulted in somewhat lower scores for the board-topper but significantly higher average survival value. Interestingly, the number of shared tops is significantly higher, almost certainly due to 2nd place being devalued in Tribute. The arrangement of shared tops appears to be a response to the devaluing of 2nd place. Alliances that survive to the end of the game are forced to be truly equitable in order to function. It is worth noting that 6 of the 13 VDL games under Tribute occurred after Weasel Moot and 3 of those ended in shared-tops, compared to just 1 of the prior 7. It appears that the shared-topping as an alliance strategy was uncovered at Weasel Moot and appeals to players under Tribute, at least for the moment.

Interestingly, games under Carnage at the Liberty Cup had on average a slightly lower draw size and a higher rate of shared board tops than under Tribute. As expected, the average topper score and survival value are both higher under Tribute than Carnage.

2.3. Gunboat

I have run three extended-deadline gunboat tournaments this year under Tribute. A fourth is currently underway using “Half-Tribute”, where the tribute rate is dropped to half a point per center (see Section 4 for a discussion of this variation). Only 14 out of 22 games have been completed at the moment, so the comparison is incomplete.

Competition	System	Games	Avg Years	Avg Draw Size	Shared Tops	Avg Topper Score	Avg Survival Value
Gunboat 1	Tribute	14	11.3	4.60	7.1%	48.2	11.4
Gunboat 2	Tribute	25	14.6	4.25	8.0%	51.0	13.0
Gunboat 3	Tribute	25	14.2	4.25	8.0%	52.8	12.2
Total 1-3		64	13.7	4.33	7.8%	51.2	12.3
Gunboat 4	Half-Tribute	14	15.0	4.33	0%	41.3	16.1

Without a clean comparison to a similar group playing under different systems, we can only draw rough conclusions from this data. There are big differences with FTF games, particularly much lower rate of shared tops. In games with lone board-tops, the average topper score and survival values are similar.

We can compare the performance under Half-Tribute, with the proviso that this sample size is small. Still, the results are in line with expectations. Half-Tribute lowers the top scores and boosts the survival value. Interestingly, the draw size is not significantly lower, despite the added elimination incentive.

3. Criticisms and responses

3.1. “Board-topping is too valuable.”

Some players have mentioned to me that board-topping seems so valuable in Tribute that allowing a lone top feels a bit like a “nuclear option”, akin to throwing a solo. While it’s difficult to quantify “nuclear”, there is no doubt that the overall board-top incentive is higher under Tribute than all other systems (see Section 1.3). The score for lone tops is usually similar to that under Sum of Squares, but because the 2nd place score is lower, the score benefit for taking the lead is much higher.

Interestingly, the worry that the allowing a board top is “nuclear” does not come up with Sum of Squares. This indicates that the perceived value of the board-top is related more to the boost in score one gets for taking the lead than the absolute value, which validates our method of measuring the board-top incentive.

So how much is too much? I think only trial and error can answer this, and it will probably depend on the context and personal preference. My sense is that the board-top incentive should be strong enough to discourage simple alliance play, but not so strong that lone tops seem too much like solos. It is certainly possible that Tribute's board-top incentive is just too strong and the increased rate of shared board-tops is a consequence. On the other hand, perhaps the demand for a shared board-top and the temptation to violate such agreements is an interesting addition to the tension between cooperation and selfishness that characterizes Diplomacy.

In any case, in Section 4 I'll propose some variations on Tribute that lower the board-top incentive.

3.2. "If you don't top, your score is terrible."

I have heard this complaint more than once and find it perplexing. It is simply a fact that Tribute awards higher scores to lower placing players than almost all other systems. The lowest score for a surviving player in Tribute is usually at least 10 points. It is almost always much less under Sum of Squares, Carnage, and C-Diplo. Draw-Size always awards higher scores to these players, and Dixie often does, but those are systems that don't or hardly differentiate between survivors at all.

In Tribute, non-topping survivors typically score 10-20 points. One could argue that scores in this range are "terrible", but I think they're better described as "mediocre". The average score for all players is always 14.3, so I think it's fair to say that scores around this value are just ok, not awful.

Whether the lower end, roughly 10 points, is substantial enough to motivate players to fight for their survival is an open question. There is anecdotal evidence from Weasels league play, VDL, and Nexus that many players do think 10 points is enough, but I have also heard some players disagree, especially since the board-topper scores so much more. The value of 10 points will depend on the context, of course, and some ways of structuring a multi-game competition will make it more valuable than others. In any case, as I discuss in Section 4 there are possible modifications to Tribute that would provide more survival incentive.

A related complaint I have heard is that 2nd place is not distinguished enough from lower places in Tribute. A player finishing just a center or two behind the board-topper but 6-8 centers above another player will only score 6-8 more points than that player, while the lion's share goes to the board-topper. This is by design, of course. A main motivation of Tribute's designers was to eliminate the notion of a "good 2nd place". We thought 2nd place should do a bit better than the lower places, but not much. We wanted to increase the board-top incentive and reduce the benefits of being a junior partner in an alliance, thereby lowering the incentive for simple alliance play. Doing this required that 2nd

place produce mediocre scores along with the lower places, which dovetailed with increasing the survival incentive, another goal.

If you think 2nd place should be worth a lot more than lower places, then Tribute is not for you.

3.3. “It’s weird.”

Tribute’s truly unique innovation is to put the board-topper’s score in direct competition with everyone else’s score. As the board-topper’s score goes up, everyone else’s score goes down in a direct way that is easy to see. This helps generate the strong balance of power incentive, which gives everyone a reason to fight the board-topper rather than each other (which may not outweigh other reasons, of course). It also means that the board-topper has a different set of incentives from the other players, and some people find this just plain weird.

The main way this shows up is with draw size. Non-toppers always score more with a smaller draw size, meaning they have some incentive to eliminate other players. However, the board-topper’s score will usually be higher with larger draw size, sometimes dramatically so. Thus the board-topper has incentive *not* to eliminate other players. Of course, this doesn’t mean the board-topper should never aim to eliminate another player; if doing so will help them gain or ensure the board-top, or even solo, then of course they should. Likewise, non-toppers may not gain as much as they think from eliminating other players (see Section 2).

Is it “weird” for a player to have incentive for something at one point in the game and then incentive against that same thing at another point? The answer depends on what you think a good strategy game should be like. If you think games should have static reward structures that don’t depend on a player’s position or the evolution of the game, then you’ll find Tribute to be problematic.

This doesn’t bother me. The only preconceived notion I have of a good strategy game is that it create difficult decisions between competing options and that success makes the decisions harder. Furthermore, as I discuss in Section 1.2, I don’t think focusing on eliminating other players is a winning strategy in Tribute. I think players should focus from the beginning on securing the board-top by whatever means necessary. If eliminating another player will increase your chances of topping the board, then definitely do that. If leaving a player alive will increase your chances of topping the board, then definitely do that. Eliminating another player for the sole purpose of increasing your score should only be pursued if the board-topper is stalemated and you have free units to do it.

It is possible that Tribute has helped us learn that topping the board and getting a high center count is actually harder in general with more survivors. If so, rewarding that is

sensible. To be clear, however, this was not an intended incentive of the designers, but rather a byproduct of Tribute's implementation of the balance of power incentive. It is possible that a strong balance of power incentive can be implemented without this "weirdness", but it is also possible that it is just novelty by another name.

3.4. "There are situations where you can lower your score by gaining a dot."

Peter McNamara advanced this objection in a letter to the editor in *Diplomacy World* #150, and I've heard it mentioned by others in casual conversation. Peter identifies a particular formal property—monotonicity—that he thinks every scoring system should have. It's meant to codify the intuition that one's score should increase with increasing center count. He presents an example where a player's score in Tribute goes down by taking a center from another player and argues that this means the system has failed the "basics".

Of the roughly 79,000 logically possible ways to arrange a player taking a dot from another, 594 result in a small loss for the taker under Tribute. These are cases where the board-topper has gotten to 14 centers or more while there are at least 5 players left and the board-topper eliminates a 1-dot power. Presumably the board-topper is trying to finish the solo at that point, sitting with 60-65 points and losing a point or two for the chance to get to 100. Note the rarity of these situations: in over 99% of dot-takings under Tribute, the dot-taker gains at least a point, often much more.

Compare with Draw-Size scoring, where only 16% of dot-takings produce an increase in score. Does this mean that Draw-Size scoring, the first scoring system, fails a fundamental test? No. Rather this demonstrates that formal requirements are poor ways to go about constructing scoring systems. A system should be evaluated on how it does in aggregate, not merely at the extremities.

As I argued in the original paper, that evaluation is best made in terms of general incentives. A debate over whether scoring systems should include a growth incentive is much more interesting than a debate over whether score should be a strictly increasing function of center count. And on a straightforward measure of a growth incentive—average change in score for taking a dot—Tribute does *better* than all other major systems (see Section 1.3). Aggregate incentive measures are simply better ways of evaluating systems than strict adherence to formal requirements.

Suppose we demanded adherence to a formal requirement like "All dot-takings should be worth at least a point". A common and well-regarded system like Carnage would fail, because 34% of dot-takings produce an increase of less than 0.01 points (normalized to 100 points per game). Has Carnage failed some fundamental principle of scoring systems? No. Rather, it means that center accumulation is not strongly incentivized on its own in Carnage. Should it be? *That* is worth discussing.

3.5. “The scores for _____ are bonkers.”

Again, the original paper endeavored to move the discussion from consideration of specific outcomes to discussion of general incentives. Every system will produce results that seem strange to some people, especially with extreme board configurations that are rarely if ever seen in practice.

Here’s an example in Tribute. A 17-11-2-1-1-1-1 board results in 83 points for the board-topper and only 11 for the second place player. The critic pounces: “Was 17 centers really more than 7 times as good as 11 centers? Absurd!”

Of course, under Draw-Size scoring, every player scores exactly the same with these center counts: 14.3. “Was 17 centers really no better than 1 center? Absurd!”

Under Carnage, the same board results in 25.0 points for the board topper and 17.8 for the 2-center power. “Was 17 centers really only 40% better than 2 centers? Absurd!”

Furthermore, suppose the game is part of a tournament, and those two players play another game, this time ending 7-6-5-5-5-1, where the 2-center power from the first game (Player B) finishes in 2nd with 6 and the 17-center power from the first game (Player A) ends with 5. Under Carnage, Player B has an aggregate score of 39.3 and player A has an aggregate score of 37.5. “7 centers over two games without even topping is better than 22 centers and a huge board top? Absurd!”

Are these results really absurd? No. Surprise at the results is due to not taking the scoring system seriously. The scoring system is part of the definition of the competition. If the scoring system values rank in a linear fashion and the players don’t play with that in mind, then that is their mistake. Likewise, if the scoring system incentivizes board-topping then players should not play in such a way as to ensure someone else tops the board, even if they get a big center count in the process.

Should we prefer systems with a strong board-top incentive? That depends on what we find important or interesting about Diplomacy. If we want to discourage simple alliance play, then yes, a strong board-top incentive is needed.

3.6. “It’s very stabby.”

In Tribute, a 1-dot stab that doesn’t lead to a board-topping position is worth 1 pt. Only 45% of possible dot-takings produce an immediate score change larger than 1 pt. By contrast, 93% of possible dot-takings in Sum of Squares are worth more than a point. Even Carnage is more “stabby” in this sense, with 66% of possible dot-takings producing more than a point (normalized to 100); this happens because there are many scenarios where one power is only one dot behind another.

In general, Tribute only provides strong incentive for stabbing if it will lead to a board top or prevent the board-topper from getting big.

Tribute does make simple alliance play riskier. 2nd place generally scores only a little better than other non-toppers, and all of those scores go down as the board-topper grows. Thus, as an alliance grows, the possible score increase from stabbing and topping the board alone gets larger and larger. A player who is not re-evaluating the risks and rewards of an alliance at every turn risks being taken advantage of by an ally who is.

Does this mean that alliances are bad under Tribute? Certainly not. Alliances are necessary to grow large enough to make a play for the board top. But they are not an *easy* strategy. If you are content to play second fiddle in an alliance, you will score roughly the same as everyone else who survives but doesn't top. Tribute is unique among major scoring systems in providing little benefit to junior partners.

3.7. "It's not relaxing." "It will lead to longer games."

Making Diplomacy more relaxing or more pleasant was not a goal of Tribute's designers. Quite the opposite. We wanted to ratchet up the intensity by providing more viable strategic options at every stage of the game. A good strategy game, in my opinion, is one that forces difficult decisions, and where success increases that difficulty. In Diplomacy, there is an essential tension between cooperation and selfishness. Because in Tribute successful cooperation leads to greater incentive for selfishness, this tension will remain more acute deeper into the game than with other systems.

This will allow for more dynamic games with more shifts of momentum and more chances to recover from early setbacks and top the board. This will undoubtedly lead to longer games than other scoring systems, as evidenced by the recent Weasel Moot.

If you want more relaxing games, Tribute is not for you.

3.8. "There's too much math."

I don't have a good response for this. In terms of computational burden, Tribute felt like an improvement over Sum of Squares because I could compute scores without a spreadsheet or calculator. But I like math; not everyone does. I wish the desired incentives could be implemented in a simpler way, but as of now I don't see how. An adjustment I propose in Section 4 is only slightly easier on the math lobe.

Of the other major systems, Carnage wins the gold medal in terms of computational simplicity. Draw-Size and C-Diplo take silver, conceptually simple but requiring a bit of math. Dixie takes bronze, more complex and requiring more math. Tribute takes last on both counts, although perhaps this will seem less of a problem with greater familiarity.

3.9. “It reduces the importance of the stalemate lines.”

I find this a really interesting complaint because it shows how people can find rather different things to like about the game of Diplomacy. I agree that Tribute diminishes the importance of the stalemate lines, because it diminishes the importance of the solo, at least relative to scoring systems like Draw-Size and Carnage. Scores between 40 and 60 are common in Tribute but rare in Draw-Size and impossible in Carnage. The solo is still quite important in Tribute, but it is truly a nuclear option in the other systems (see 3.1). Stalemate lines are important because they provide a significant barrier to getting a solo, one that can figure heavily in negotiations and strategic thinking in the late game.

I have had many discussions with David Hood (who runs DixieCon and developed its proprietary scoring system) about scoring, and for him the stalemate lines are a central feature of the game of Diplomacy. In his view, the extra barrier they create is a crucial hurdle that requires long-term planning and deep cunning to surmount, characteristic of the challenges unique to Diplomacy.

I think this is a perfectly valid way to view the game. It’s just not how I prefer to view it. I see Diplomacy as a competition for power requiring a balance of cooperation and selfishness. So it’s natural that I would value board-tops, which identify the player who accumulated the most power. And it’s natural that David would not value mere board-tops, which have not overcome what he sees as the central challenge of the game.

Neither David nor I are right or wrong on this. Rather, we are describing different games. The scoring systems, as extensions of the primary Diplomacy ruleset, define those different games.

4. Possible adjustments

4.1. Surrenders

One of the unpleasantnesses of Draw-Size scoring is that games can get dragged out for the sole purpose of eliminating a player. To combat this, a variant was created in which players can surrender, or agree to be left out of the draw even though they still control supply centers. A player will surrender if they know their position can’t be defended against a coordinated attack by the remaining players, and they aren’t able to throw a solo to any of them (or those players agree to forgo the solo opportunity). Obviously this leads to shorter games. A system that does not allow surrender votes is known as DIAS (“draw includes all survivors”), and one that does is referred to as non-DIAS.

Some players take a “cold, dead hands” approach to non-DIAS systems and refuse to surrender under any circumstances. I have not yet played at a tournament where surrender votes are allowed, but I suspect I’d fall into this category. If I’ve gone to the

trouble of traveling somewhere to play Diplomacy, then I will hold on to every last unit and force someone to actually pass up a solo, not just say they will. But again, I look to Diplomacy for intense strategic challenges, not relaxing experiences. I recognize not everyone feels this way.

Surrender votes can be incorporated into Tribute in a simple way. Those with centers who don't participate in the draw still receive 1 pt for each center they control but do not get a survival bonus and do not pay tribute. Note that this will lower the board-topper's score under most circumstances.

While this could only be determined by trying, I suspect that allowing surrenders would play out differently in Tribute than in Draw-Size and Dixie. In Tribute, the small powers have a way of hurting the other non-toppers short of throwing a solo, giving them more leverage. In any case, I think you wouldn't see many instances of a 4-5 center power surrendering, which is not uncommon in Draw-Size and Dixie.

4.2.Changing the tribute rate and/or threshold

If we wish to lower the board-top incentive (see Section 3.1) or raise the survival incentive (see Section 3.2), there are some simple adjustments available to us. The two most straightforward are to raise the threshold for paying tribute or to lower the tribute rate. However, raising the tribute threshold would significantly reduce the value of smaller board tops such as those with 7 and 8 centers. Because these are fairly common in time-limited games, it's best to leave the threshold at 6.

Here are two workable variations:

Half-Tribute: Reduce the rate of Tribute to 1/2 point per supply center. This lowers the board-top incentive considerably while raising the survival incentive, but also damages the balance of power incentive.

Semi-Tribute (or Half-Tribute-Plus): Reduce the rate of Tribute to 1/2 point per supply center and add a 6 point bonus for topping the board. In other words, survivors get 1 point per center + $60/N$, where N is the number of survivors. The topper gets 6 points plus, for each center the topper has over 6, 1/2 point from each of the other survivors. This has the added benefit of simplifying the survival bonus calculation, since 60 is divisible by 2, 3, 4, 5, and 6.

Here are the incentive measures for these options, compared with other systems:

Incentive measures for Tribute variants

System	Growth	Survival	Board-Top	Dominance	Balance of Power
Carnage	1.44	0.00	3.53	0.00	-15%
C-Diplo	4.01	1.01	20.49	1.00	-2%
Dixie	1.03	11.90	0.68	0.95	-9%
Draw-Size	0.53	16.65	0.00	0.39	-12%
Sum of Squares	4.35	0.26	5.97	5.65	11%
Tribute	4.65	5.82	21.55	5.42	27%
Half-Tribute	3.06	8.89	11.28	3.42	10%
Semi-Tribute	3.50	7.89	15.81	3.39	12%

Both Half-Tribute and Semi-Tribute boost the survival incentive significantly and drop the board-top incentive.

Do they boost the survival incentive enough to convince players that survival is worth fighting for? Was Tribute's value already high enough? Only time and experimentation will tell. Dixie's value of 11.90 is clearly high enough, and Sum of Squares's value of 0.26 is clearly too low. "Good enough" must be somewhere in between, although that may vary somewhat with context.

And what is an optimal value for the board-top incentive? Sum of Squares (5.97) is clearly too low (for my tastes), but is Half-Tribute's 11.28 high enough? And is Tribute's 21.55 really too high? I suspect Semi-Tribute's 15.81 is high enough to discourage simple alliance play.

The balance of power incentive is harder to gauge. Sum of Squares has a higher value (11%) than other major systems, but that is almost entirely due to the 2nd and 3rd place powers. Smaller powers have little to lose and little to gain by getting a dot, regardless of who they take it from. Tribute and its derivatives spread the incentive more evenly over all players. Full Tribute definitely promotes balance of power play in practice, and I wouldn't want to sacrifice that. Experimentation will be required to see if Half-Tribute or Semi-Tribute still have the same effect. As I discuss in Section 1.3, it's possible that the board-top and dominance incentives together are better measures of balance of power incentive than the next-dot measure I've been using.

For those wanting to drop the board-topper's score a bit while maintaining the unique features of Tribute, I'd recommend trying Semi-Tribute first. I suspect Half-Tribute's board-top incentive is a bit too low, especially for smaller board-tops.

5. Conclusion

Tribute has shaken up the Diplomacy landscape, prompting deep discussion about styles of play as well as new efforts at even more radical system design. This can only be good for the hobby.

Whether Tribute has lasting impact or not remains to be seen. All the theory about incentives won't matter if people don't like to play it more than other systems. The early returns have been largely positive, but there are some vocal curmudgeons. To them I'd offer this: Tribute may indeed be the worst scoring system, except for all the others.