Use Chances Created, not Assists or Expected Assists: Predicting playmakers' future creativity

Three ways to measure playmaker creativity

Assists, Expected Assists, and Chances Created (A, xA, and CC) are three standard ways to measure a playmaker's creativity. Rather than using them to admire past performance (the perspective of the fan), our focus is on predicting future performance (the perspective of the club, the player, the agency, and so on). And when it comes to prediction, we argue CC should be preferred over both A and xA as the way to measure *playmaking*. We use the following table to frame the initial discussion. Then in the next section we present empirical evidence that that favours CC.

	Assists	Expected Assists	Chances Created
Freeze the frame	✗	✓	✓
Objectivity	✓	✗	✓
Numerosity	✗	✓	✓
Simplicity	✓	✗	✓

Freeze the frame (when the passer's contribution is over). Every assist must start with a chance created. But once the ball has left the passer's boot, there is nothing he can do to influence whether his unique contribution turns into an assist or not. The shooter's touch has to be good enough to get a decent shot away to evade the defence and the goalkeeper. We can write the elements involved in an assist schematically as:

$$A = CC + (touch + shot\ quality + defence + goalkeeper + luck)$$

About 1 shot in 10 results in a goal, so it is no wonder the contribution of the passer is usually a minority stake in the goal itself. Yet when a goal results, full credit is given to the passer; and no credit when no goal is scored, even for a tap-in missed. By not freezing the frame at the moment the passer's contribution is finished, all those other contaminating factors (bracketed in the above 'equation') are allowed in as noise into *the assist* as a measure of playmaking.

Both xA and CC freeze the frame (at least in theory) when the passer's contribution is finished. CC adds 1 for every chance created, irrespective of the probability of a goal resulting; whereas xA replaces that 1 with a number between 0 and 1, depending on the probability of a goal resulting. Also, xA gives credit to the playmaker where no shot was actually taken (hence CC would not be incremented), say if the shooter spills the ball in a promising position. Thus, xA seems like a reasonable upgrade over CC, in that it weights better chances more than lesser ones.

Objectivity. However, xA relies on hundreds of coders to assess those probabilities, and wherever subjective interpretation occurs, so does the potential for human bias. *Hindsight bias* is the tendency for people to judge an event as more likely to have happened if they already know that it has actually happened, as when the commentator shouts: “You just knew he was going to score!” No-one is immune to hindsight bias: the financial analyst who knew a crisis was brewing (after the event); the physician who feels sure they would have chosen the correct diagnosis (once they know what the illness is); and the xA coder who overestimates the probability of a goal attempt being scored (because they know that it *was* scored) – and contrariwise underestimates the probability of goal attempts that they know *were not* scored. The upshot is that xA may not be as objective as it should be, because knowledge of whether a goal was actually scored or not (*i.e.*, knowledge of A) bleeds back into the assessment of xA. And as we have already argued, A is a very noise-ridden when evaluating the passer. Later, we present empirical data to establish these claims.

Numerosity. “The fewer goals there are in a sport, the more impact random events have.” (Rasmus Ankersen: Ted Talk, Football analytics Brentford-Midtjylland). Assists are even less numerous than goals (every assist implies a goal, but not *vice versa*), so Ankersen’s claim could be made of assists too. Both are examples of the *Law of Large Numbers*. Because there about 10 chances created for every 1 that turns into an assist, CC has *numerosity* in a way that A lacks. This fact alone should ensure that CC is a much more accurate measure of a player’s ability to create chances than A is of a player’s ability to make assists. The following example should either clarify or confuse,

Coefficient of Variation (CV) is a unitless measure of dispersion or spread, defined as standard deviation / mean. Measures with less dispersion (more precision) are better. A Poisson process has a mean = variance = L. So, $CV(L) = L^{-½}$. Assuming both A and CC are Poisson with means of a and c, respectively; and that $c \approx 9a$ (to keep calculations simple), then $CV(a) \approx 3CV(c)$. In conclusion, because of its greater numerosity, CC will be a much less dispersed, therefore more precise, measure than A. Note that these calculations parallel the translation of standard deviation into standard error by the $N^{-½}$ multiplier. These working assumptions of a Poisson do not need to be exactly correct for the general conclusion still to hold.

Finally, according to Opta, every completed pass generates an xA, even if no shot / header results. So it too has numerosity.

Simplicity. CC and A are both easy to understand. This makes both CC and A intelligible to the football world; but it also means that inter-coder reliability should be very high – everyone knows when a chance has been created, or an assist made (including coders). Of course, there will always be some contentious instances, but they are the minority. By contrast, few understand exactly how an xA is estimated (and data providers like to keep it that way), so the rest of the world has to take it as a matter of faith that someone, somewhere, must know what they are doing.

The data favours CC.

Because a player's assists are contaminated with significant contributions from other players, all outside his control, and because assists occur too infrequently to get a reliable fix on performance, it is not likely to be a measure that predicts well into that player's future performance. Later, this will be confirmed empirically. But first we concentrate on the shortcomings of xA.

Bias. xA is biased. If an archer aims at the target, but the arrows land consistently to one side of the centre, this would be bias. Similarly, if a coder (or coders) consistently over-estimate the real value of xA, that would also be bias, and this is exactly what happens. The definition of xA suggests that in the long run xA ≈ A. In practical terms, if we add up all xA for a season, it should be roughly the same as the sum of all A. The following table shows that this is not the case.

	Open play		**Set piece**	
	A	**xA**	**A**	**xA**
English1	529	598	156	86
English2	708	878	229	146
English3	725	863	250	147
English3a	710	801	216	122
English4	658	854	262	147
French1	491	538	96	60
German1	549	693	143	78
Italian1	561	623	138	82
Spanish1	462	535	143	85
All	5393	6383	1633	953

English1, English2 refer to English first and second tiers, 2025-26, etc. English 3a is data for 2024-25.

Over the 9 leagues, xA overestimates by about 1000 in open play, but underestimates by about 700 in set piece play. These are an order of magnitude greater than the differences we might expect due to random fluctuations, and are consistently problematic across all leagues. We conclude that coders show large amounts of bias when judging xA. This problem would not be insurmountable if we could be sure that inter-coder bias was reliable. Then, as overestimations of xA would be consistent, we could deflate every xA estimate by the necessary percentage. For instance, if xA = 0.10 resulted from a set piece, using the 'All' data in the above table we would recalibrate it to:

xA = 0.100 x 1633 / 953 = 0.171.

However, the next analysis pulls this rug from underneath xA. The workaround assumes that xA is reliably coded, but it is not.

Test-retest analysis. By chance, we have a legacy dataset from the previous season for EFL League 1 (identified as E3a). This allows us to model the accuracy which last year's A, xA, and CC for a particular player will predict the same player's A, xA, and CC in this year, and predicting future performance goes to the heart of the matter. 240 players were identified who had played in League

One for both 2024-25 and 2025-26 seasons, and who had sufficient minutes (>= 450 mins in each year separately). The correlation matrix between all 6 measures (A, xA, CC) for two years is shown below. Data were standardised to per minute data. All the revealing information is in the bottom left quadrant (3 x 3) as it contains the correlations of last year's measures with this year's.

		this season (2025-26)			previous season (2024-25)		
		A	**xA**	**CC**	**A**	**xA**	**CC**
this season 2025-26	**A**	1					
	xA	0.565	1				
	CC	0.400	0.726	1			
previous season 2024-25	**A**	0.193	0.185	0.336	1		
	xA	0.309	0.427	0.497	0.520	1	
	CC	0.393	0.494	0.589	0.391	0.755	1

First, the highlighted leading diagonal contains the test-retest correlations for A, xA, and CC (per 90 mins) between seasons 2024-25 and 2025-26. Since CC is highest (R = 0.589 versus 0.193 and 0.427), someone's CC per 90 will carry over from one year to another better than their A or xA will. This is to be expected for A, because Assists are a noisy measure of talent, due to the playmaker having no control over whether his CCs will be converted to goals, and thereby assists. Consequently, the test-retest correlation is low for A (0.193). But now we see that xA is also a noisy measure of talent: here, the noise comes from coder unreliability. We conclude that CC is more reliable than A or xA. This is all good and well, but we may really want to predict A and xA rather than CC. After all, assists imply goals, and goals win games.

However (and remarkably), last year's CC predicts not only this year's CC, but also predicts this year's A better than A itself or than xA (because 0.393 > 0.193 and 0.309). Similarly, xA is better predicted by CC (because 0.494 > 0.185 and 0.427). We can also reverse time's arrow, and show that CC (2025-26) better infers past A, xA, and CC (2024-25) than either of the other 2025-26 variables (because 0.336 > 0.193 and 0.185; and 0.497 > 0.427 and 0.309).

We can come at the same issue test-retest predictivity from a different angle by running three multiple regression analyses (OLS) with 2025-26 season's A, xA, and CC in turn as the y variable, and 2024-25 measures of A, xA, and CC as the x variables. Without overburdening with too much detail, just the t-values for the coefficients are in the following table ($t_{critical}$ = 1.97, df=236, 2-tail). In all three analyses, CC is the only significant variable. This allows us to conclude that not only is CC most predictive, but neither A or xA significantly enhance predictivity over and above CC. All that can be captured *is* captured by CC.

		predicting 2025-26 measures		
		A	xA	CC
using	A	0.65	-0.76	1.76
2024-25	xA	0.04	1.63	0.77
measures	**CC**	**4.09**	**4.67**	**6.28**

Finally, in the correlation matrix, the two values highlighted in green suggest that xA does a much better job of predicting a player's A if they are both *within* the current year. This is exactly the profile we expect of hindsight bias, in which knowledge of A seeps into xA, and becomes a spurious boost to predictivity *within* the same year. But when it comes to predicting *between* years, which is the acid test, that component of A that has seeped into xA is now no longer insider information but irrelevant noise. Like the pupil who has peeked at the answers to an exam: it may help if they are the answers to this year's exam, but not if they're last year's answers.

In conclusion, CC does a better job than either A or xA in predicting a player's future creativity (whether that future creativity is to be measured by A, xA, or CC). Because A includes so much that is nothing to do with the playmaker, we cannot see how it ever could become a reliable predictor. On the other hand, xA does do a better job than A, but its undoubted potential is undermined by having to involve human judgement with all its biases and errors incompletely corrected. Without an extensive clean-up of xA by data providers, chance creation is left as the most reliable measure to predict future CC (of course), but also future A and xA.

Drawing conclusions from a single dataset is a limitation, but when season 2026-27 is finished these questions could be revisited, armed with a more extensive set of data. Unfortunately, without access to either coder behaviour, or the background statistics and patterns that underpin the estimation of expected assists, xA itself will remain a black box.

John Doyle, Football Analytica
21-07-2026