Are 3D Printed Dice Fair?

U. Hermann
Udo's 3D World, Germany

(Dated: 21 October 2025)

We investigate the statistical fairness of 3D-printed dice produced by fused-deposition modeling (FDM). Using an automated rolling rig and computer-vision readout over 7 000 – 21 000 throws per sample, we evaluate deviations from uniform probability via Pearson's χ^2 test with five degrees of freedom.

Dice printed with 100% infill and corrected geometry show no significant bias within 95% confidence limits, while intentionally modified or flawed prints, such as cavities, asymmetric infill, or "elephant's foot" distortion, exhibit measurable skewness. A simplified game-level simulation reveals that the strongest bias increases win probability from about 50% to 60%. Our results define practical printing criteria for unbiased dice and demonstrate that fairness testing is feasible with consumer-grade equipment.

Introduction

"Alea iacta est"

The rapid advance in 3D printing made it very easy to develop new board games, since any part of the game (except for cards and maybe the board) can be produced at minimal cost. Many board games use dice to add in randomness and tension, and the combination with 3D printing allows for a cheap and easy production of fully customized dice for new board games.

When presenting 3D printed dice to other players, their fairness is questioned immediately. 3D printing is not perfect and introduces many artefacts (and also offers some rather unethical options) that can influence the fairness of a die. The aim of this work is to determine how big the influence of these artefacts is to conclude whether dice made with a 3D printer are fair or not.

Methods

We are using 3D printed dice made from a Bambu Lab A1 3D printer[1]. The focus on FDM is mainly due to the availability of multi-coloured prints, that would allow to take dice directly out of the printer without having to perform extra steps like painting the indented numbers. To automatically roll dice, a machine is introduced that rolls the dice and takes images of them, using AI image recognition with a convolutional neural network to classify the results.

The mechanical part consists of a chamber with stairs that is flipped upside down by a small servo motor to achieve a random rotation of the die. The die then falls back down onto the transparent glass and is rotated towards the camera to show the bottom face. This process is repeated until the desired number of images was taken.

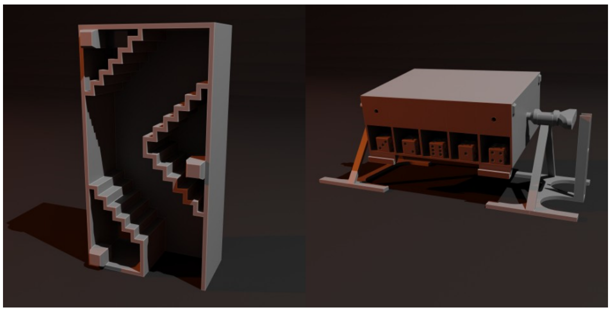

FIG. 1. 3D render of the experimental apparatus. The left part shows the inner structure to make the die roll, the right part shows the full machine as it was used.

To determine the fairness of a die, Pearson's χ^2 test[2] was performed, calculating the χ^2 value for every die. With 6 faces, a die has 5 degrees of freedom and the required χ^2 value to say with a certain probability that the die is not fair is:

Probability	90%	95%	97.5%	99%	99.9%
Required χ^2	9.236	11.070	12.833	15.086	20.515

TABLE I. Critical values of Pearson's χ^2 statistic for a six-sided die (five degrees of freedom) at different confidence levels.

For our tests, we decided to go with a 95% probability, so the threshold to say if a die is fair or not is at χ^2 of 11.070.

With 10 000 rolls being more than a normal die usually withstands (Let's assume a game of Settlers of Catan, with an average of 71 turns per game[3], played every month, over 10 years equals 8520 dice rolls), it might be that the die is worn out during the measurement, which could cause rising or falling χ^2 values. Therefore, the χ^2 value is calculated for sequences that are only 1/10 of the whole measurement, to evaluate whether any trend is happening (continuously wearing out the dice), or a sudden change is observed (parts of the die breaking off).

Since the χ^2 value still means that 5% of all dice are above that value, the final evaluation metric was decided to be the average of 10 χ^2 values, each performed on subsequent parts of equal size of the whole measurement set. This removed high spikes from the measurements and allowed a much better evaluation whether a die is fair.

Test Objects

The machine holds 5 different dice (but is easily extendable). For our test, these 10 dice were used

Die 1: The “normal one” – a 12 mm white die with black dots.

Die 2: The printed one – a 3D printed 16 mm grey die (100% infill) with indented, painted, red dots.
Upper face during 3D printing: 6.

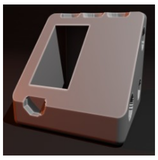

Die 3: The cheater’s die – like die 2, but with a large cavity towards the “3” face. No visual difference towards die number 2 can be seen.
Upper face during 3D printing: 6.

Die 4: The bad die – like die 2, but with an exaggerated “Elephant’s foot”, a bottom side (the “4” face) that is widened by 1 mm. An “Elephant’s foot” is a very common 3D printing artefact.
Upper face during 3D printing: 3.

Die 5: The weird one – a skew die[4] that has tilted faces, but the creator claims for the die to be fair, even backed by a scientific paper.
Upper face during 3D printing: 6.

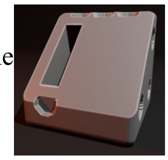

Die 6: The cheater’s die 2 – the same as die 3, but with a smaller cavity to halve the displacement of the center of gravity compared to die 3.
Upper face during 3D printing: 6.

Die 7: Lightweight 6 – die 2, but with deeper extrusion holes on the 6 side, in the hope to see (or not to see?) an effect.
Upper face during 3D printing: 6.

Die 8: The not so bad one – die 4, but with a much smaller elephant’s foot, 0.3 mm.
Upper face during 3D printing: 3.

Die 9: The hollow die – the same as die 2, but without a 100% infill, to see if a non-filled die causes problems (e.g. due to top and bottom faces not being equally heavy).
Upper face during 3D printing: 3.

Die 10: The corner – A die with one corner sticking out, as it happens sometimes to the lowest corner in resin printed dice. The corner is at the intersection of the 3, 5 and 6 faces.
Upper face during 3D printing: 6.

Results

A total of 7 000 – 21 000 measurements were performed and evaluated. The evaluation was performed in Python and the result was put together in an image.

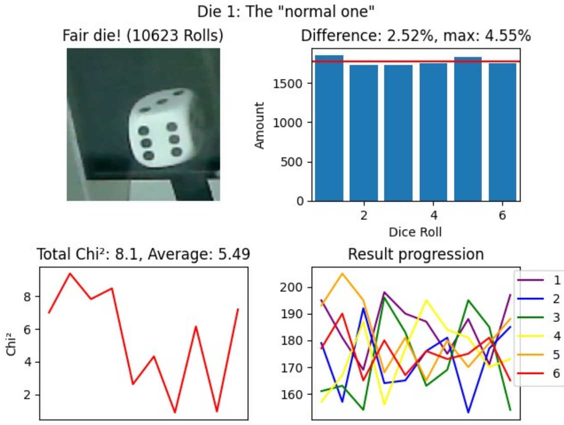

FIG. 2. Die 1: The “normal one”.

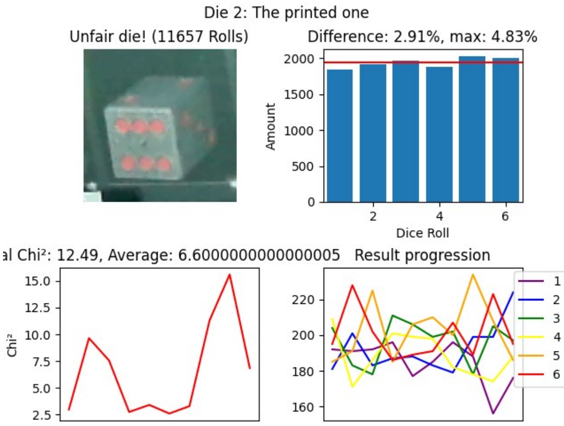

FIG. 3. Die 2: The printed one.

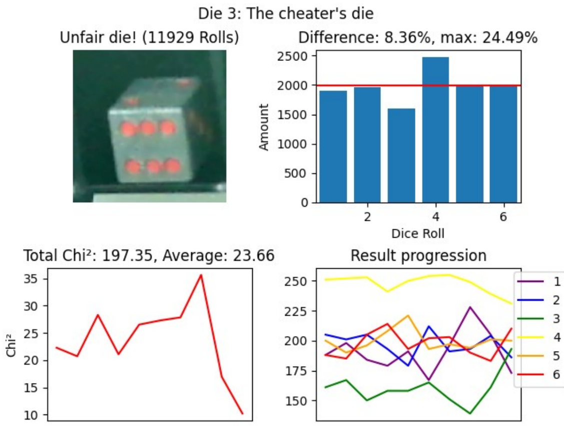

FIG. 4. Die 3: The cheater’s die.

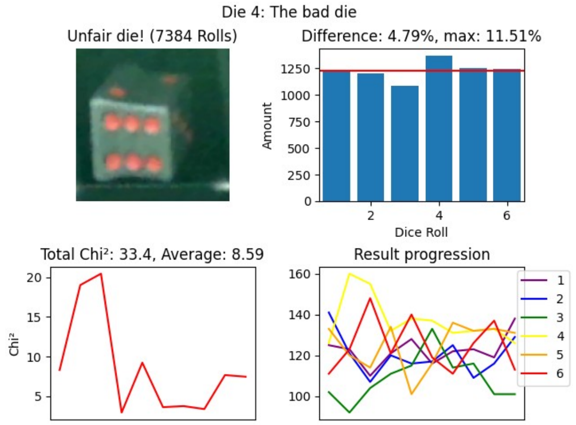

FIG. 5. Die 4: The bad die.

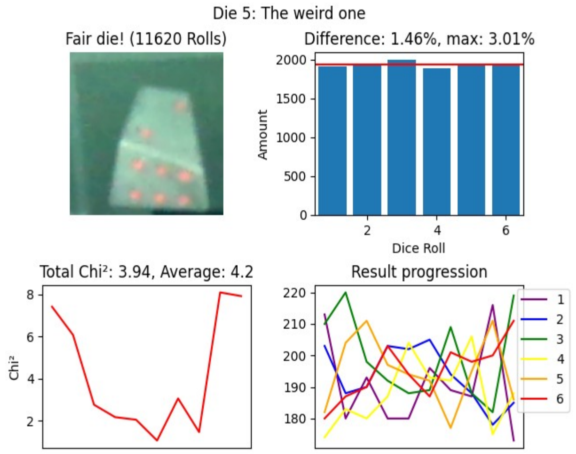

FIG. 6. Die 5: The weird one.

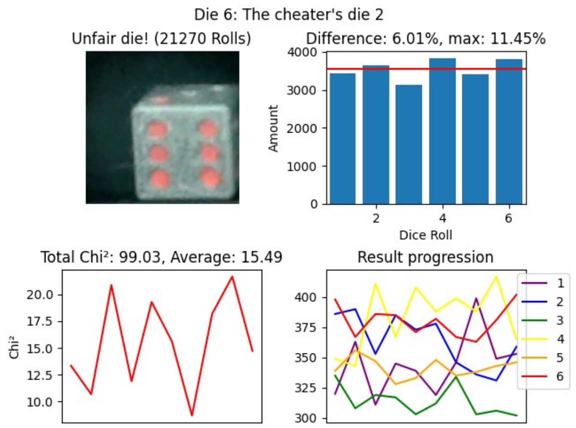

FIG. 7. Die 6: The cheater's die 2.

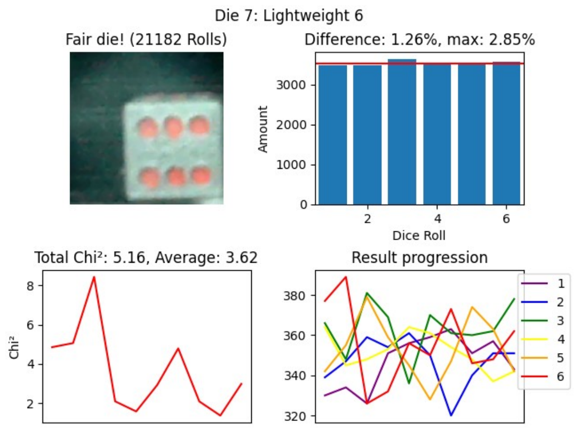

FIG. 8. Die 7: Lightweight 6.

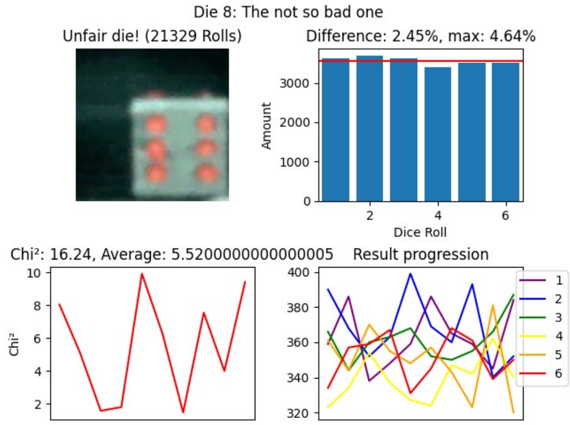

FIG. 9. Die 8: The not so bad one.

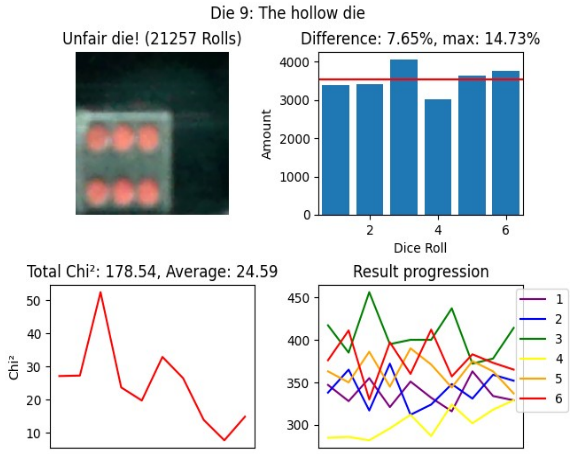

FIG. 10. Die 9: The hollow die.

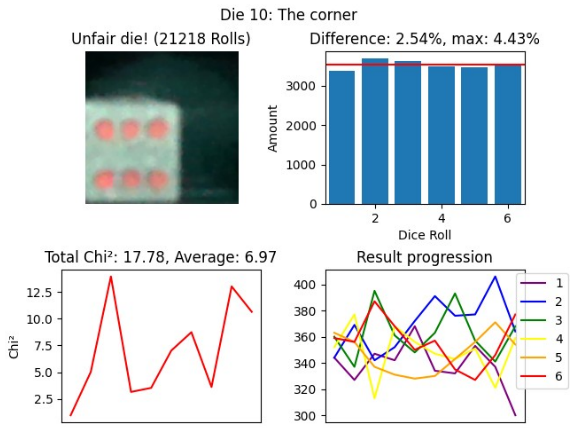

FIG. 11. Die 10: The corner.

The influence on a game

With all these results, one big question is still open: How big is the influence on a real tabletop game? To test this, a simplified game of Grimdark Future[5] was played: Two squads of 10 Battle Brothers each without extra equipment were used. A Battle Brother has a single attack hitting on 3+, and defends against these attacks because of AP(1) on 4+. If only half of the squad is there, a morale test has to be performed on a 3+, followed by another roll of 4+ if it failed. If the unit fails both of these, it counts as wavered and will not attack in the next round. This was performed 100 000 times with different percentages of cheating. To simulate the cheater's die, for every roll of 1, a chance was introduced that this die is a 6 instead. The chance taken for this was the maximum difference from the die results, being 24% for the cheater's die and 11% for the cheater's die #2.

Die error	Chance of winning
0% (Perfect die)	50.3%
5% (Normal/Printed die)	52.3%
11% (Cheater's die #2)	54.8%
24% (Cheater's die #1)	60.0%

TABLE II. Simulated win probability in a simplified tabletop scenario as a function of die bias. Each value represents the mean probability of victory over 100 000 simulated battles using modified dice with the indicated level of bias.

Chances of winning in relation to the die error.

The cavity in the cheater's die was nearly fully expanded to the sides, so there is not too much extra space to "optimize" this die. So with a roughly maxed out cheater's die, the chances of winning only rise to 60%, being an increase of roughly 20%.

Since this method of cheating moves the center of gravity of the die, throwing the die into a liquid with a higher density (e.g. water infused with sugar or salt) will show a face that always turns up at the top, so this method can be used to check for weighted dice.

Results discussion

It can be seen that dice 1, 5, and 7 show fair behaviour. This is especially surprising for die 5, though having such tilted faces, but as stated in the sources, this die shall be fair and this experiment finally proves that! The effects caused by the artefacts of die 7 seem small enough to not have an influence on the result.

Dice 3 and 6, as expected, shows an unfair behaviour. Die 6 shows a much smaller effect than die 3, but it is still visible. Die 9 actually also belongs into this category – the automatic slicing added 3 full layers at the bottom of the die, but 9 full layers at the top, also resulting in an uneven distribution. Since face 3 was on top, this "heaviest" face lands on the bottom most, where our camera can see it.

Die 2 is labelled as unfair. Looking at the diagrams, no clear preferred side can be seen and it seems that due to statistical noise, this die was labelled as unfair. More measurements are needed to verify this, and since only the pre-last measurement bin shows an extreme increase in the χ^2 value, this can be regarded as a random effect.

This observation poses a big threat for 3D printed dice - do I really know whether my opponent has fair dice? With less than 100 rolls, the difference vanishes below statistical noise, so this cannot be determined quickly, but requires a dice rolling machine as presented here. A solution to this might be the floating method: Throwing dice into water and seeing which side comes on top, even after disturbing the die several times. This method might be more useful than rolling the die 10 000 times, but might still be impractical for a tournament. Another solution for this might be that dice are shared between players, so both players benefit from the manipulated results, or the tournament provides fair dice to players.

Looking at die 2, it can be seen that on the 6 face, more material is removed than on the 1 side. This might also have an effect, so another experiment was started with die 7, to increase this effect. All the holes on the 6 face were made twice as deep, with no observable effect, leading to the conclusion that this small amount of mass can be neglected when talking about fair dice.

Die 4 shows that a precise geometry is important for a fair die. An Elephant's foot of 1 mm is extremely exaggerated, but small values like 0.2 to 0.4 mm might occur, depending on how good the 3D printer is calibrated. Performing a measurement afterwards and ensuring that the bottom of the die is not wider than the top might be a suitable way to produce fair dice. Increasing the rounded edge's radius might also compensate for this behaviour. Die 8 has a more realistic elephant's foot, with only 0.3 mm, and the results show that this does affect the fairness, but much less than die 4.

Die 10 also shows the importance of a precise geometry, having a slight unfair behaviour.

Conclusion

3D printed dice can be fair! A precise geometric shape needs to be ensured and especially an infill value of 100% needs to be set during printing. If these things are considered and no active manipulation of the die is occurring, it is possible to 3D print dice that are fair and can be used in games.

[1] Bambu Lab A1 3D Printer, https://eu.store.bambulab.com/
[2] Pearson's chi-squared test, https://en.wikipedia.org/wiki/Pearson%27s_chi-squared_test
[3] Board Game Breakdown: Settlers of Catan, the basics, https://www.alexcates.com/post/board-game-breakdown-settlers-of-catan-the-basics
[4] Skew Dice by Carlos Luna, https://www.printables.com/en/model/616349-skew-dice
[5] Grimdark Future by OnePageRules, https://www.onepagerules.com/games/grimdark-future