Can One Hear the Fate of a Coin? [draft!]

Kartik Chandra[*]

October 4, 2020

"Even a single coin has music in it."

— Tom Stoppard's *Rosencrantz & Guildenstern are Dead*

(after Rosencrantz' coin lands "heads" several dozen times in a row)

Are coin tosses fair — that is, do tossed coins have a 50:50 probability of coming up "heads" or "tails"? Diaconis, Holmes, and Montgomery [1] say *no.* A coin toss is a physical process, and given a distribution over the parameters (how fast it spins about its diameter, ω, and how long it spins, t), we can infer a distribution over result (heads, tails), which will not necessarily be uniform.

The trouble is measuring these parameters. Coin flips are fast. With some practice you and a friend can use a stopwatch to clock a typical toss to be around 0.5 seconds. Then, you can count the number of flips by attaching a small ribbon (or dental floss) to the coin and counting the number of twists after a toss. But this would be extremely tedious work.[1] You might hope to save time by using a video-camera to record a coin flip and counting the flips, but typical cameras sample images at not more than 60 Hz, which is far too slow. High-speed cameras can capture at higher frame rates, but only in short bursts: capturing usable sequences, and then using image analysis to extract useful data, is still a challenge.

Why not use sound, which is typically sampled at a luxurious 44.1 *k*Hz? If we could "hear" the angle a coin makes with respect to the microphone at a given point in time, then we could work out ω by "listening" to a slowed-down recording of a coin-tosser.

It turns out that you can indeed "hear" this angle. When you toss a coin by striking it with your thumbnail, it starts to resonate and ring (this is the familiar audible "ding!" sound). The resonance can be decomposed into a linear combination of several vibrational modes, which depend on the geometry and material properties of the coin. Each vibrational mode has a characteristic frequency. For a 2016-D US quarter, audible vibrational modes occur at 8.7 kHz, 9.3 kHz, and 15.4 kHz. These values are intrinsic to the coin and can be measured empirically or computed using simulation software [3].[2]

An important characteristic of resonance is that it does not necessarily radiate acoustic energy equally in all spatial directions. Just like a loudspeaker sounds louder from some directions than from others, a resonating coin has an *anisotropic acoustic transfer function* [2], which may differ across vibrational modes. Thus, the subset of frequencies that an observer hears most prominently depends on an observer's relative angle to a resonating coin. As the coin rotates, these frequencies weave in and out at a rate proportional to ω. Figure 1 shows two potential modes ("trampoline" and "potato chip") and the dominant directions of their acoustic transfer functions ("perpendicular" and "coplanar," respectively).

[*]kach — at — cs.stanford.edu

[1]Another technique uses stroboscopic illumination to track a coin with one side painted a different color.

[2]In fact, you can detect counterfeit coins by comparing their resonant frequencies to those of real coins.

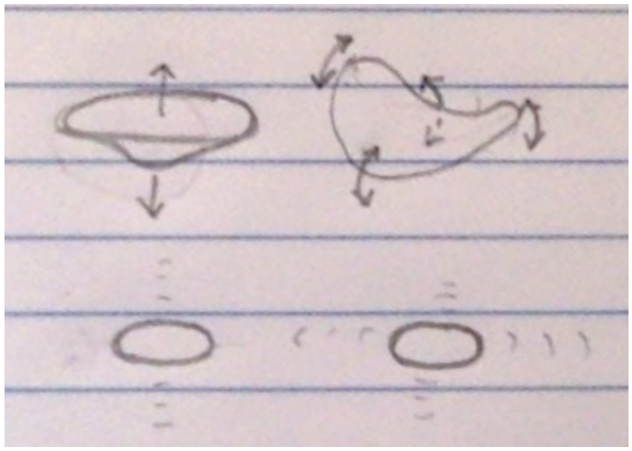

Figure 1: Schematic of two possible vibrational modes of a coin and the dominant direction of their acoustic transfer function.

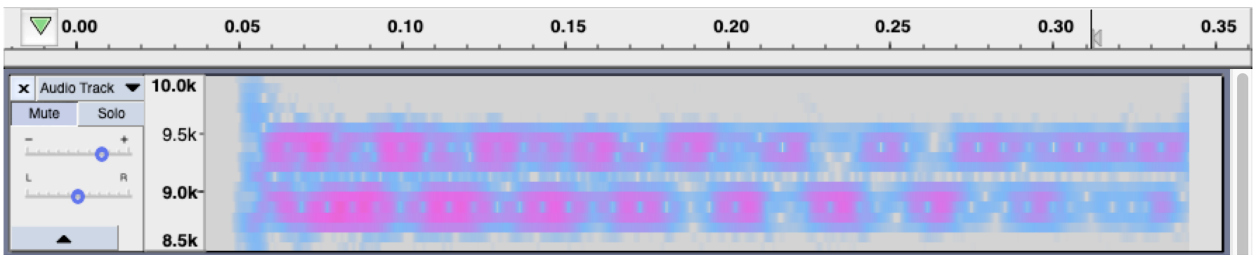

Figure 2: Spectrogram of a toss of a 2016-D US quarter, focusing on the 8.5-10 kHz band.

In general, computing acoustic transfer functions is a tricky simulation problem. However, for the purpose of measuring a coin's ω, the exact acoustic transfer function need not be computed. If we measure *any* periodic variation in amplitude at a given frequency, we can determine ω up to a small whole-number constant c, which counts how many "lobes" the transfer function has (in the case of Figure 1, $c = 2$). Because earlier experiments have already determined rough bounds on the distribution of ω, we can use our data to infer c.[3]

Figure 2 shows a spectrogram of a toss of a 2016-D US quarter, produced using the Audacity software package. The axis along the top shows time in seconds; the scale on the left shows frequency in kHz, and the color represents intensity scaled such that pink is the "loudest."[4] Notice that there are clearly two phase-shifted sinusoidal signals present, at frequencies roughly corresponding to the two lower resonant frequencies of a US quarter. Now, at each frequency, we can simply count peaks to determine (up to a factor of c) how fast the coin is spinning.

How can we automate this kind of analysis? Figure 3 shows the intensity (i.e. pink-ness in Figure 2) of a signal at 9.3 kHz over time (normalized). As expected, it looks like the graph of $|\ \sin(\omega c \cdot t)\ |$; the absolute value is because the signal from the "far side" of the coin takes over after each half-turn. To take a meaningful Fourier transform of this signal, we can square it, which yields a pure sinusoid from which a frequency can be extracted.

[3] We also know that c must be even because of the heads-tails symmetry of the coin.

[4] The spectrogram was computed using a Hanning window of size 512. The specific size is not too important, as long as it is short enough to capture the envelope frequency and long enough to distinguish the two carrier frequencies.

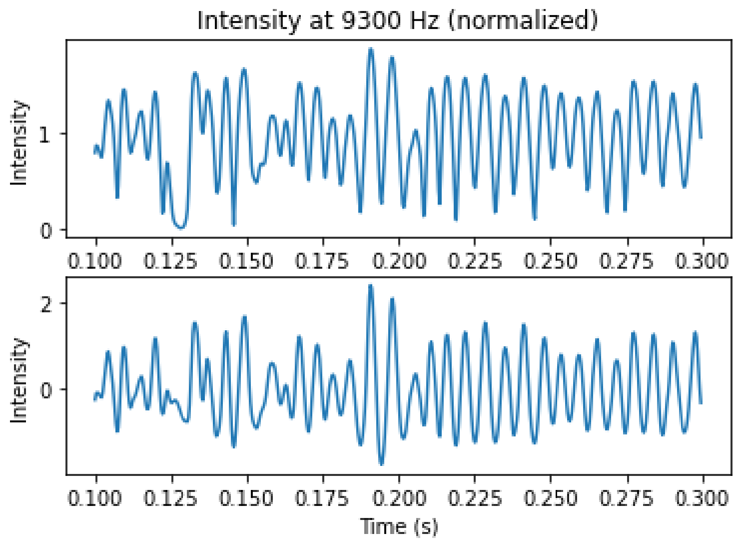

Figure 3: The normalized frequency profile over time (top) is the absolute value of a sinusoid; once squared and centered (bottom), it is a pure sinusoid.

Experiments & Results

This method is extremely practical and can be carried out at home by anyone. No special hardware was needed to record the audio: it was captured with an iPhone 7 in a room with plenty of ambient background noise.[5] Data from several tosses can easily and automatically be extracted from a single long audio clip with simple heuristics. The entire processing software is around 30 lines of code.

Figure 4 shows the results for ≈ 200 "vigorous" coin tosses collected over a 15-minute interval, which is potentially the largest coin-dynamics sample size ever collected.[6] Based on this distribution, it seems reasonable to conclude that $c = 2$, which would bring these results to approximately the same order of magnitude of the quoted frequencies of "36-40" in [1]. We conclude that (the author's) tosses of US quarters spin at $\boxed{55 \pm 10 \text{ Hz}}$.

Of course, the results presented here have much higher variance than the 36-40 range, and, assuming $c = 2$, they skew higher. This might be explained by two factors. First, [1] uses half-dollar coins, which weigh twice as much as quarters and therefore naïvely should spin at half the rate when struck with the same energy. Second, this author is a clumsy amateur coin tosser who took several hours to even *find* a physical coin in his house — by no means as consistent as a trained magician. Perhaps, however, this situation is more representative of the "average American coin-tosser," which would suggest that the results in [1] require some adjustment to apply to daily life.

References

[1] Persi Diaconis, Susan Holmes, and Richard Montgomery. "Dynamical bias in the coin toss". In: *SIAM review* 49.2 (2007), pp. 211–235.

[5] Background noise is typically low-frequency and does not affect the 8-10 kHz band.
[6] Audio recording available on request.

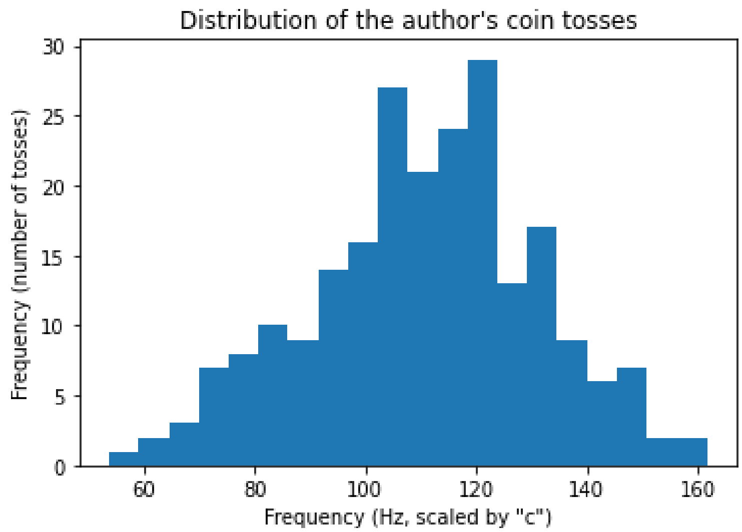

Figure 4: Empirical results for ≈ 200 coin tosses. From this plot we can attempt to infer c, and therefore the distribution of ω.

[2] Doug L. James. "Physically Based Sound for Computer Animation and Virtual Environments". In: *ACM SIGGRAPH 2016 Courses*. SIGGRAPH '16. Anaheim, California: Association for Computing Machinery, 2016. ISBN: 9781450342896. DOI: 10.1145/2897826.2927375. URL: https://doi.org/10.1145/2897826.2927375.

[3] Emerson Steed. "Coin identification through natural frequency analysis". In: *Brigham Young University. https://me363. byu. edu/sites/me363. byu. edu/files/Emerson_Steed_CoinIdentification. pdf* (2015).