The Physics of the Torpedo Bat: An Update

Alan M. Nathan*

Department of Physics, University of Illinois, Urbana, IL 61801

(Dated: April 29, 2025)

I. INTRODUCTION

Now that the initial excitement about torpedo bats has died down somewhat, I have begun to think more generally about how baseball bats actually work. Actually, I first started thinking these thoughts way back in 1998, when I had just begun a sabbatical year and was looking for interesting diversions from my usual research. Over the next few years, I learned a lot about bats, how they work, what makes one bat better than another, etc. That initial work resulted in my very first **academic paper** on the physics of baseball.

In view of the recent interest in bats resulting from the torpedo bat, I have returned to thinking about these things once again. I wanted to focus in particular on the idea that the mass of the bat in the vicinity of the impact point (the "local mass") plays an important role in the performance of the bat. The essential idea seems fairly obvious: The more mass there is behind the collision, the harder the ball will be hit, with the usual "all other things equal" caveat. And in fact, that idea is true. However, the reason *why* it is true is less obvious and more subtle than most people appreciate. The purpose of this brief article is to explain the physics behind this idea. All the results shown here are *calculations*, using computer codes that have previously been tested successfully against experimental data.

II. A TOY MODEL

To fully understand the physics of the baseball-bat collision requires some fairly high-level mathematics, as you might appreciate if you simply glance at my **paper** on the subject. It is important to our understanding to isolate the fundamental physics involved once all of the

* a-nathan@illinois.edu

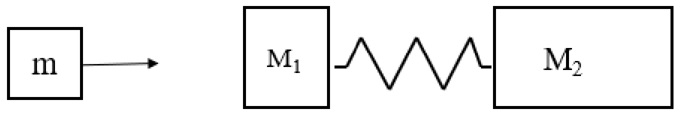

FIG. 1: Toy model for the ball-bat collision. The ball is a block of mass m, while the bat consists of two blocks of mass M_1 and M_2 connected by a spring. The ball is incident from the left.

mathematical details are stripped away. Therefore I have created a highly simplified "toy model" of the ball-bat collision that allows us to clearly see the important physics involved. The model is depicted in Fig. 1. The bat is modeled as two blocks of mass M_1 and M_2 which are connected by a spring, while the ball is modeled as a single block of mass m. The ball is incident on the bat, which is initially at rest, with initial speed v_i. It subsequently bounces off the bat with rebound speed v_f. The quantity of interest is the collision efficiency q:

$$q = \frac{v_f}{v_i} \, . \tag{1}$$

In fact this type of experiment is routinely done in laboratories where a baseball is fired at high speed out of an air cannon onto a bat at rest and the very same q is measured. It is one of the metrics of bat performance. In real life, there is energy dissipation in the ball, as measured by its coefficient of restitution, or COR. The dissipation in the ball is not essential to the point I want to make, so for simplicity I will assume no dissipation (i.e., COR=1).

If our model bat were a single block of mass M without the spring, then the result for q can be derived using only freshman physics:

$$q = \frac{1-r}{1+r} \qquad r = \frac{m}{M} \qquad \text{single block.} \tag{2}$$

In this simple picture, q depends only on the ratio r of the ball to bat mass, which controls how much the bat recoils from the collision. In the limit $r \ll 1$, as in a ping pong ball colliding with a bowling ball, the ping pong ball bounces back with its initial speed (i.e., q=1). In the opposite limit $r \gg 1$, as in a bowling ball colliding with a ping pong ball, the bowling ball just keeps on going in its initial direction and the same speed (i.e., $q = -1$).

Now let's consider the more complicated problem with the bat being two blocks connected by a spring. Such a toy bat can both recoil, as the one-block model, and also vibrate. This is a bit more complicated problem to solve than the one-mass problem, but it is nevertheless fairly straightforward to a relationship for q analytically. In doing so, we find that q depends not only on the masses but also on one additional parameter which is the product of the vibrational frequency f of the masses+spring system that constitutes our bat and the collision time T. The frequency f is determined both by the masses of the two blocks and the "stiffness" of the spring. For given masses, a greater stiffness means a higher frequency; lesser stiffness means a lower frequency.

This problem can be solved exactly. Rather than show you a formula relating q to the masses and fT, it is more useful to show a graph of an example, as in Fig. 2. In this example, the ball has a mass of 5 oz and the collision time is fixed at T=1 ms. In the red case, M_1 and M_2 are 10 oz and 20 oz respectively. In the blue case, they are 12 oz and 18 oz, respectively. In both cases the total mass is 30 oz. The stiffness of the spring is varied so that fT is in the range 0 to 2.5 (note that fT is dimensionless). We see that the resulting q does indeed depend on fT in a way that looks a bit complicated. Despite this, there are two extreme cases that have a very simple physical interpretation.

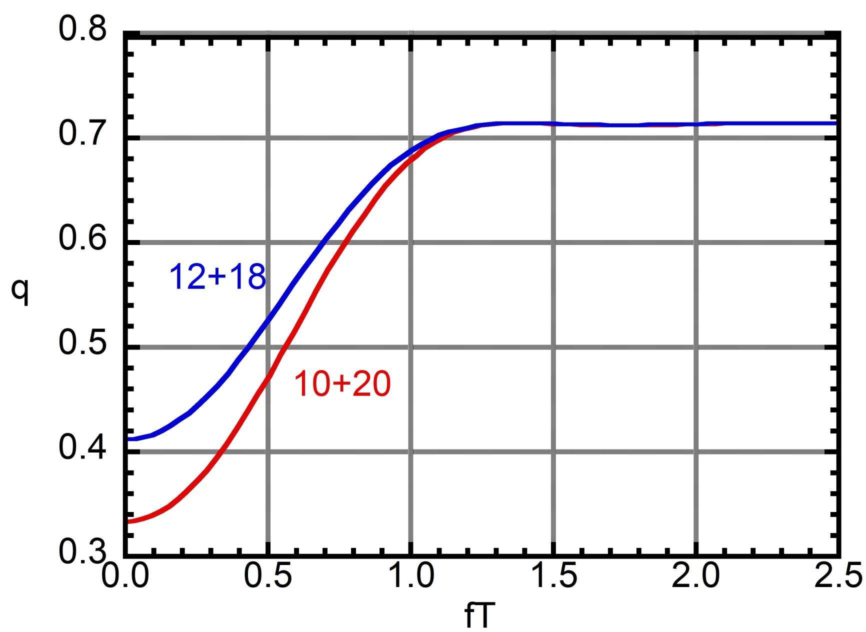

FIG. 2: The collision efficiency q plotted versus the factor fT, where f is the natural frequency of the bat and T is the collision time. The red curve is for the case M_1=10 and M_2=20; the blue curve is for the case M_1=12 and M_2=18.

For $fT \geq 1.5$, corresponding to a very stiff bat spring (like the coil springs in a car), the ball pushes slowly on the bat, which recoils as a whole without compressing the spring. The resulting q is exactly the same as in the single block case with $r = 5/30$ for both the red and blue bats. Another way to say that is that on the long time scale of the collision, the bat is completely rigid and the ball "sees" only the total mass and not the individual pieces that make up that total mass. In that limit, the red and blue bats perform *identically* despite the different sharing of the total mass between the two blocks. Note that since f is just the inverse of the period of oscillation of the spring, the limit $fT \geq 1.5$ applies whenever the collision time is comparable to or longer than the period of the spring.

Now consider the opposite extreme, $fT \ll 1$, corresponding to a very loose spring, sort of like a slinky. In that case, the ball bounces off M_1 quickly before M_2 has time to react. M_1 eventually recoils backwards, compressing the spring and starting M_2 in motion. The resulting motion of the bat is a combination of translational recoil and vibration. On the short time scale of the collision, the two pieces of the bat are essentially uncoupled from each other and the ball only sees M_1. Indeed q is exactly what one would expect from the single block model with mass M_1, so that $r = 5/10$ and $5/12$ for the red and blue bats, respectively. In this case, rearranging the masses made a substantial difference in the performance of the two bats. This very simple example confirms the popular notion that the larger the local mass, which is M_1 in this example, the harder the ball can be hit (i.e., larger q). Finally note that the condition $fT \ll 1$ is equivalent to the collision time being very small compared to the period of the spring.

In the general case, one can still interpret the results as the ball colliding with a single block with an effective mass M_{eff} intermediate between M_1 and $M_1 + M_2$. The full calculation is needed to find M_{eff}, or equivalently q, and the recoiling bat has some mixture of translation and vibration that depends on fT.

Before closing out this section and turning to an examination of realistic bats, I want to emphasize the main thing we learn from the toy model:

- **The notion that a larger local mass results in a higher collision efficiency can only be understood by treating the toy bat as a dynamic system, capable of both translation and vibration.**

III. RELATING BATS TO TOY MODEL

So what does this toy model have to do with the real world of bats? That's a fair question. And it has an answer. A bat seems like a perfectly rigid body. But it is not. It can vibrate. You can easily induce these vibrations yourself by holding a bat lightly between two fingers and tapping it lightly. Depending on where you tap the bat, you can hear the vibrations and even feel them in your fingers. Take a look at this **video** and you will see what I mean.

A bat is essentially a nonuniform cylinder, meaning that the diameter varies along its length. One way to model this bat is to imagine that this cylinder is sliced up into much thinner cylinders, say half an inch thick, where each of these slices is connected to its two neighbors by a spring.

With this picture, you can sort of see how it relates to the toy model. The ball strikes one of the slices, starting that slice in motion. This ends up stretching the springs connecting it to its neighbors, thereby starting those slices in motion. And on it goes. In effect, when the ball strikes the bat, a little localized wave is created at the impact point, which then propagates down the bat in each direction and bounces back and forth as it reflects off each end. And just like in the toy model, the effective mass that the ball "sees" depends critically on the contact time. For very long impact times compared with the time it takes the wave to propagate the full length of the bat, the ball sees the entire bat, just as in the toy model. For more typical impact times and for impacts in the barrel, the time it takes for the wave to propagate to the knob, reflect, and return to the impact point is longer than the ball-bat collision time, so the mass at the knob end of the bat cannot possibly affect the collision. In fact, that is the origin of "the grip doesn't matter" **discussion**. But the wave also bounces off the barrel end of the bat, and with the shorter distance it reaches the impact point while ball and bat are still in contact. So the effective mass of the bat during the collision is a complicated function of the actual distribution of mass, the pulse propagation time, and the impact time.

To see how all this plays out to affect the collision efficiency is not so simple. As you might imagine, keeping track of all this motion is not so easy to do by hand. But it is very easy to do with a computer. And that is exactly how my code for simulating the ball-bat collision works.

IV. COMPARING TWO BATS

Now let's see how these ideas apply to two realistic bats and a realistic ball-bat collision. Consider the two bats whose diameter profiles are shown in Fig. 3. The standard bat has the usual profile, with the maximum barrel diameter, 2.5", occuring right at the barrel tip. The torpedo bat modifies the standard bat by transferring mass from the tip toward the handle, so that the maximum diameter, also 2.5", is near the sweet spot, approximately 7" from the tip. For distances further from the tip than 12", the bats are identical. The resulting properties (mass, length, balance point, and moment of inertia) of the two bats are very similar and not likely to result in any differences in performance. The simulation consisted of firing a ball at each bat, initially at rest, with an impact speed varying between 145 and 170 mph, depending on the impact point. The net outcome is to find the collision efficiency q. All of the necessary physics and mathematics needed to carry out this computation can be found **here**. The results are shown in Fig. 4.

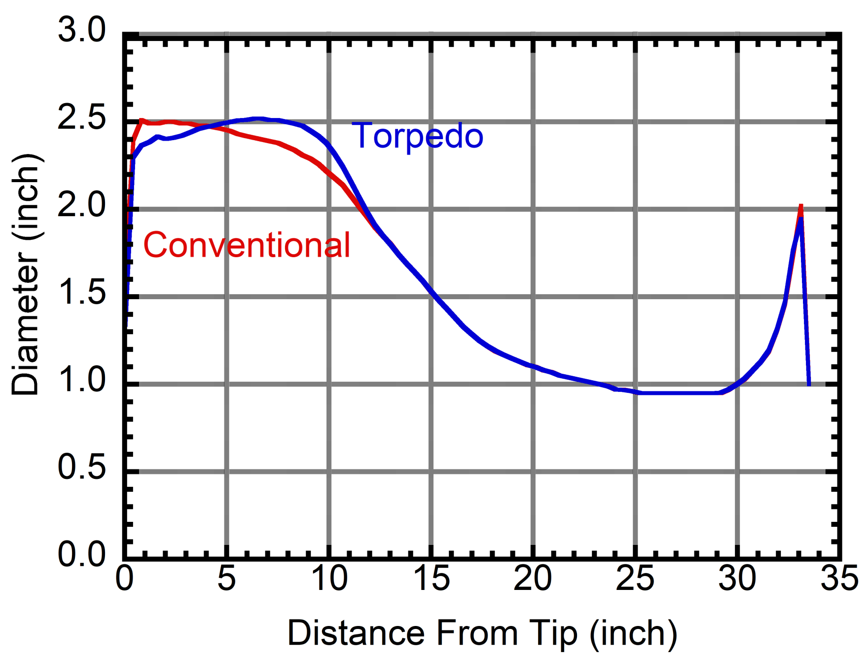

FIG. 3: Diameter profile of a standard (red) and torpedo bat (blue).

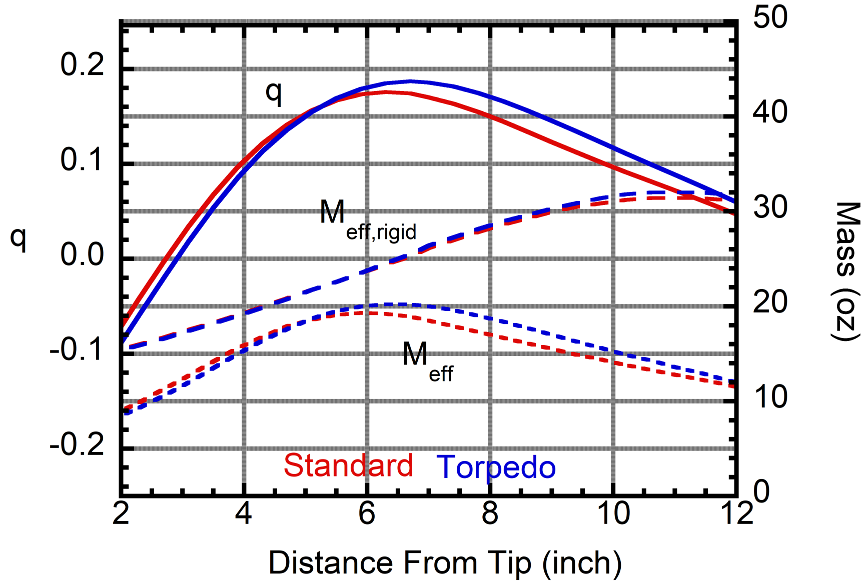

FIG. 4: Results of a realistic collision between a baseball and a standard bat (red) and a torpedo bat (blue), plotted as a function of impact location along the barrel. q is the collision efficiency, M_{eff} is the equivalent effective mass, and $M_{eff,rigid}$ is what the effective mass would be if the bats were perfectly rigid or if the collision time were very long. For the latter, the two bats perform nearly identical.

The collision efficiency q varies with impact location in the usual way, with the maximum slightly further from the barrel tip for the torpedo bat as a result of its different mass distribution. Comparing Figs. 3 and 4, we see that, roughly speaking, q is larger for the standard bat at locations where the local mass (i.e., the mass of the bat at the impact point) of the standard bat is larger than that of the torpedo bat (i.e., closer to the tip than the sweet spot) and smaller where the local mass is smaller (i.e., farther from the tip than the sweet spot). This confirms the popular notion that a larger local mass will result in a harder hit ball.

As a side note having nothing to do with the primary purpose of this article, Fig. 4 shows that the amount by which the collision efficiency of the torpedo bat exceeds that of the standard bat inside the sweet spot is greater than the amount by which the opposite is true outside the sweet spot. As a result, the range of q above some specified (but arbitrary) minimum is greater for the torpedo than for the standard. In somewhat loose parlance, that means the sweet spot of the torpedo is greater than that of the standard.

But how we got there is important. As with the toy model, one can interpret q as the collision of the ball with an effective mass M_{eff}, which is also plotted in the figure. And as in the toy model, the effective mass has a simple interpretation: It is the equivalent mass that results in the same q and that recoils with a motion consisting of a mixture of translation, rotation (see footnote below), and vibration. We see that q and M_{eff} essentially track together as the impact location is varied. The effective mass is related to the actual local mass but not in a simple way. Nevertheless, the bat with the larger local mass will have both a larger q and a larger M_{eff}.

To emphasize the importance of bat vibrations, I have also plotted $M_{eff,rigid}$, which is what the effective mass would be if the bat were perfectly rigid (i.e., no vibrations)[1] or if the collision time were very long. The rigid-body effective mass is equivalent to $M_1 + M_2$ in the toy model and essentially coincides for the two bats, meaning that these bats would perform identically if vibrations were neglected, just like the large fT limit in the toy model. The effect of vibrations can be seen as the difference between $M_{eff,rigid}$ and M_{eff}, which always results in the former exceeding the latter. That difference is minimized right at the sweet spot, which is the impact location where bat vibrations are minimized because the local stiffness of the bat is maximized. The differences in effective mass of the two bats is due to their different vibrational properties, which in turn is due to their different mass distributions in the barrel. It is precisely those vibrational differences that lead to the differences in performance.

V. CONCLUSION

- **When comparing the performance of two bats, it is necessary to treat the bats as dynamic objects, capable of vibrating, rather than static rigid objects.**

VI. ACKNOWLEDGEMENTS

It is a pleasure to acknowledge very useful discussions with Prof. Patrick Dufour and to thank him for a critical reading of this article. The underpinnings of this work were developed long ago and benefitted greatly from many lively discussions with Prof. Rod Cross.

[1] The reason for the variation of the rigid-body effective mass with impact location is that, unlike the toy model, a real bat can recoil both by translational and rotational motion about the center of mass. The rigid-body effective mass is largest for impacts at the center of mass, for which only translational recoil is possible.