


Newton's Cradle.

by Donald Simanek




The physics toy and physics demo sold as "Newton's cradle" is also called "colliding balls", "Newton's spheres", "counting balls", "impact balls", "ball-chain", the "executive pacifier", and even, believe it or not, "Newton's balls."

This document makes no claim to be a complete treatment of this device. The extensive literature references should be consulted by those who wish to delve deeper. We will attempt to shed a little light on things not often mentioned in elementary textbook treatments, and suggest some experiments one might do to test certain assumptions about the physics of elastic and inelastic collisions.

1. The apparatus.

The apparatus usually consists of an odd number (5 or 7 is common) of identical steel balls each suspended by a bifilar suspension from a sturdy frame. The balls are carefully aligned along a horizontal row, just touching each other.



If you don't have one of these and haven't seen one in action, please play a bit with this [interactive flash animation](#) demonstrating the idealized behavior of this apparatus.

Animated image provided by [Raven Black](#), and used here by permission of the artist.

When the ball on one end is pulled aside and allowed to swing as a pendulum, it hits the next ball. But the outcome is fascinating, the one ball on the far end is knocked away from the others with the same speed as the first ball had initially and all of the other balls remain nearly at rest. If you pull back two balls and let them strike the others, two balls are ejected from the other end, and all the other balls remain nearly at rest. Why does this happen? Why are these the only outcomes that occur? Why not others?

We shall refer to this as the 'standard behavior' and the standard observed outcome, for the purposes of discussion. We are quite aware that this is an idealized outcome, and that the real apparatus doesn't quite achieve it, though it comes quite close. We are also aware that deviations from the ideal conditions (differences in mass, spherical vs. cylindrical masses, some balls touching, some not) can cause very interesting deviations from expected behavior and are severe tests for any model of system behavior. Some of these are discussed in the references at the end of this document.

Why has this become a standard demonstration in physics courses? What important principle is it supposed to show? Usually it is "advertised" as a demonstration that momentum is conserved in collisions. Well, the outcome certainly does illustrate that, but so does every other mechanical interaction you might care to consider, whether it be elastic or inelastic. This particular apparatus, cleverly designed to be nearly elastic, is a *special case*, and the full generality of conservation of momentum is not demonstrated by it. Some books say that this demo shows that **both** energy and momentum are conserved in a collision. That's closer to the mark. But still, this apparatus is a special case: collisions between spherical nearly perfectly elastic balls of equal mass, size and composition. And those special conditions are responsible for the intriguing and special behavior we observe.


Unfortunately this raises in the inquiring student mind all sorts of questions, such as "What if the balls were of different size, mass, shape or composition?" And that opens a whole can of worms that could lead one far astray and consume a lot of class time. It can easily lead into discussion of energy-dispersive systems and the importance of impedance matching. Seeing the demonstration raises a very sticky question: "How do the balls "know" that if you have N balls initially moving, that exactly N balls should swing out from the other end?" This is the feature of this apparatus that justifies the name "counting balls", for the system seems to "remember the number" of balls that were pulled aside initially. This is the big question that the elementary accounts do not answer satisfactorily. Answering all of these questions would be fine for an upper-level university physics course, but are hardly suitable for a freshman or high school course.

How do those little balls "know" how many must be ejected at the other end?

Sometimes textbooks suggest a simpler version of the apparatus: marbles rolling on a grooved

track. This presents even stickier problems, for the result depends on conservation of energy (including linear and rotational kinetic energy terms), conservation of momentum and conservation of angular momentum. Still worse, friction, rolling resistance, slipping on the track and momentum exchange with the track during the collision all affect the outcome.

Do not expect to find all of the answers in this web document. Consult the journal references at the bottom of this document if you want to delve deeper. As a point of departure we begin by limiting our discussion to the classic apparatus: identical perfectly elastic spherical balls. Only later will we consider systems with various shapes, masses and materials.





Homebuilt Newton's cradle using ceramic cabinet knobs and Meccano/Erector parts.
Stereo for cross-eyed viewing.

2. Inadequate textbook treatments.

Textbooks and internet sites often tell lies about this demonstration. Here's a few of these lies, with my comments in square brackets:

1. The observed outcome is the **only** one that conserves both energy and momentum. [Not so. There are others.]
2. The balls aren't really touching, so a series of independent two-ball collisions occurs. [This is not a necessary assumption. Many commercial versions of this apparatus do have a slight gap between the balls. That improves the desired performance somewhat, but the surprising outcome is essentially the same whether or not they are touching.] And that raises the question





Homebuilt adjustable apparatus using "hi-bounce" balls, map tacks string, small (#20) binder clips, binder posts Erector parts and magnets to hold it to the blackboard.

"Why does the successive collisions model predict the right answers even when the conditions for it are not met?"

3. The observed outcome is due to the finite speed of sound (an elastic compressional wave traveling down the line of balls). [This is also not a necessary assumption, for the standard outcome would be seen even if the speed of sound were infinite! However if the speed of sound were very slow relative to the ball velocities, we might have different outcomes.]
4. This demo requires wave equations and solution of 2N simultaneous equations to understand. [This may not be necessary. The materials of freshman physics may be sufficient to answer **some** of the obvious questions. Such an analysis could reveal some very basic and important physics at that level. The more general and mathematical analysis may obscure some interesting conceptual details.]

Some fake explanations try to convince the reader by citing a particular example:

before: $V \rightarrow$ balls at rest
 $0 \quad 00$

after: $00 \quad V \rightarrow$
 $\text{at rest} \quad 0$

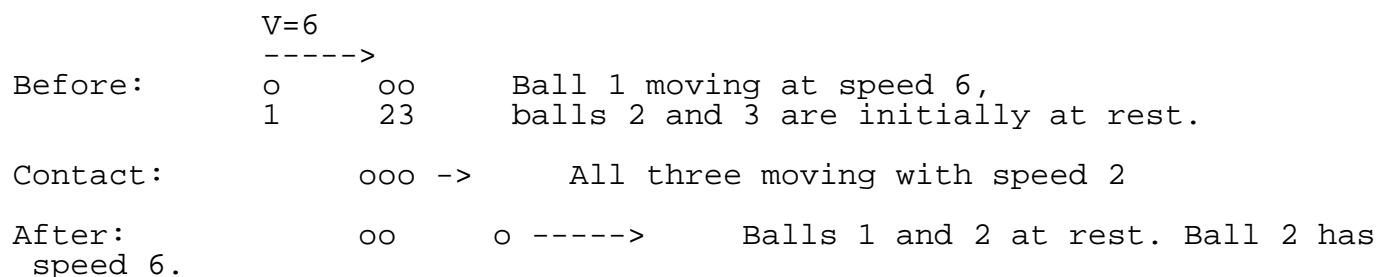
Consider three balls. Balls 2 and 3 are stationary. Ball 1 hits ball 2 with speed V . Ball 3 moves away with speed V leaving balls 1 and 2 stationary. [Yes, that's what happens in the real world.] Momentum and energy are both conserved.

Why are other results impossible? Consider the hypothetical outcome: Ball 2 and 3

move off with speed $V/2$, leaving ball 1 stationary. Answer: This conserves momentum, but not kinetic energy.

The math is easily checked, and every statement is true. But this not the end of the story by any means, nor is it a proof, nor does it give any insight into the problem. You do not prove the impossibility of *all* other outcomes by showing that *one* other outcome is impossible. That's an elementary fallacy of logic.

Are there other outcomes that satisfy conservation of energy and momentum, yet are **not** observed to happen? Yes, and they are easily found. For simplicity, take all masses to be 1. Ball 1 has initial speed V , balls 2 and 3 are initially at rest, touching each other.



This is the observed outcome, which we will call "case 1", summarized below. But why does case 2 not occur?

Case 1, the observed outcome

Quantity	Initial 1	Initial 2	Initial 3	Final 1	Final 2	Final 3
velocity	6	0	0	0	0	6
momentum	6	0	0	0	0	6
kinetic energy	18	0	0	0	0	18

Case 2, two balls emerge at the same speed.

Quantity	Initial 1	Initial 2	Initial 3	Final 1	Final 2	Final 3	Net
velocity	6	0	0	-2	4	4	
momentum	6	0	0	-2	4	4	6
kinetic energy	18	0	0	2	8	8	18

Case 2 clearly satisfies conservation of energy and momentum. Yet it is not observed to happen. Suspicion focuses on the processes happening during the impact, when the three balls were in contact for a brief time interval. What's going on there? The balls deform elastically near the point of impact.

Can you find any other hypothetical situations that would conserve energy and momentum but do not happen?

If we can answer this question for the three-ball case we might gain insight into the N-ball general case. In fact, if we look carefully at the two-ball case we might learn something about how the elastic properties of the balls store and release energy.

Some textbooks and web sites tell you nothing other than a description of the behavior of the system and note that this behavior satisfies the conservation of energy and momentum. Perhaps they are "playing it safe" by not attempting to answer the obvious questions that "inquiring minds want to know."

So, with these trivial distractions out of the way, what is the explanation? Why is just one (of many) momentum and energy conserving outcomes selected by the laws of physics, to be the only outcome that happens?

Discussion

The final velocities after collision depend not only on the initial energy and momentum of the balls, but on the nature of the impact. The impact causes deformation of the balls at the points of contact and the temporary storage of energy as elastic energy. In the case of three balls, there are two contact interfaces. During a brief time interval compression occurs at both interfaces, and some of the initial energy is stored at each one. During this time interval all the balls are moving with nearly the same velocity, and conservation of momentum determines that velocity exactly.

Then separation occurs after the elastic energy is released back to the balls. Since there are two interfaces, each interface stores only a part of the total compressional energy. In the case of two equal mass balls with ball 2 initially at rest the observed outcome tells us that just enough energy is released at the contact interface to cause ball 1 to slow it to a stop. From Newton's law we know that there's equal size and oppositely directed forces at the contact point acting on balls 1 and 2 at all times. But the elastic force acting leftward on ball 1 is acting opposite to its speed, succeeding only in slowing that ball to a stop. Since the force acting on ball 1 is opposite to its displacement, the work done by that force is negative. That means that ball 1 is doing positive work on ball 2, increasing the speed of ball 2. Ball 1 loses kinetic energy and ball 2 gains the same amount of kinetic energy.

In the system of three identical elastic balls, the elastic forces at the interface of ball 2 and 3 are also oppositely directed during the time all three are in contact, moving together. The force on ball 2 slows it to rest. The force on ball 3 increases its speed. So ball 2 and 3 must separate. This is why the hypothetical outcome in case 2 (above) does not occur in the perfectly elastic collision of three identical balls.

Other energy and momentum conserving outcomes do not occur in the three ball case, for more obvious reasons. If ball 1 impacted stationary balls 2 and 3 at speed V , an outcome with ball 1 emerging at speed V with balls 2 and 3 still at rest would conserve energy and

momentum. But it would mean that ball 1 never hit balls 2 and 3. Or, perhaps it passed right through them!

Other considerations include the question of the speed of the compression impulse through the metal of the balls. Some of these issues are discussed in the bibliography references.

Perfect elasticity?

Nothing in nature is perfectly elastic. Physics textbooks often say "assume the process is perfectly elastic". What they mean is "Assume that energy losses due to inelasticity have negligible effect on the outcome". We know that this toy isn't perfectly elastic for several easily observed reasons. (1) We hear noise when the impact(s) occur, and that is due to energy being lost from the system as sound energy. (2) As successive collisions occur, the system eventually comes to rest. So the kinetic energy of motion must have either been stored, or dissipated as heat and sound.

We will, in this discussion, assume that the collision is perfectly elastic, that is, the sum of system kinetic and potential energies is conserved. We need not fuss about the fact that collisions are never perfectly elastic. They can be close enough to perfect for this analysis to help us understand what's going on. This is an example of simplifying the problem as a method for gaining insight, then relaxing the simplification later if necessary.

The word "elastic" has two distinct meanings in mechanics. A "perfectly elastic material" is one that can be deformed and then return to its original condition exactly, without energy loss. A "perfectly elastic collision" is one that involves contact and deformation of two bodies for a short time. After the collision the energy of the two body system is the same as before the collision. **But** a collision between two perfectly elastic bodies may not be an elastic collision. Consider a bell, made of highly elastic steel struck by a clapper of similar material. The materials are both very elastic. But after the collision of clapper and bell, considerable energy remains in the bell, vibrational energy, which is slowly dissipated as radiated sound energy and some is dissipated as thermal energy, heating the material of the bell. The materials are very elastic, but the collision is not, for after it's all over, a considerable fraction of the mechanical energy is gone through these two processes of dissipation. The clapper separated from the bell before the stored elastic energy in the bell could be converted to kinetic energy.

Idle question: What would be the behavior of a Newton's cradle made of identical suspended bells (without clappers)? What would it sound like?



Newton's Bells?

Newton's oddball cradles.

We remarked that some versions of this apparatus have one larger ball at one end. This version reveals another remarkable result. If the large ball is at the far end, and the first ball is pulled back and released, then *only the first and last balls move after the collision*, all the others remain nearly at rest. Does this work in reverse? Pull back the large ball and release it. After collision all the balls except the ones at the ends remain at rest. There's a symmetry here, which may be a clue. The first and last balls behave as if the others only transmitted the energy and momentum without being disturbed.

What if we had a row of an odd number of balls, all identical except the one in the middle, which is heavier than the others? Surprisingly, the cradle still works in the usual way. Well, this is not so surprising when you think about it.

Now try a string of balls with the heavier ball second from the end. This disrupts the simple behavior we have seen before.

These experiments tell us a lot about what's going on, which we wouldn't have learned from the usual (all equal) arrangement of balls.

Does the initial speed matter?

Using the standard Newton's cradle with equal elastic balls, experiment with different initial speeds. Pull back one ball just a bit and release it. The one at the far end moves away, as if all the balls in between played no role. Pull back one ball a large distance, and the one at the other end moves away a large distance, as if all the balls in between played no role. This is telling us something important. It seems that all that determines the outcome is the impulse from the initial impact, so long as it is of short enough duration. The balls in between seem to be only transmitting the impulse from one end of the row to the other.

This sheds light on the results of the oddball cradle. The only effect that the initially moving ball has is to produce elastic compression at the first interface for a very brief time. Does the rest of the system "know" exactly what produced that compression? The mass of the initially moving ball didn't compromise the fact that all balls except the end ones ended up at rest.

So we are motivated to test this, by substituting some other kind of impact, perhaps by hitting the ball at one end with a small hammer. Try as you might, you cannot do this to duplicate the results of one ball hitting another. Why? One reason is that it's nearly impossible to duplicate an impact of such short duration as that of the collision of two balls. You end up pushing the whole line of balls forward for a longer time. To duplicate the standard behavior of the cradle the initial impact between the first two balls must be over before the impulse is transmitted to the

other balls in the row.

Modeling Newton's Cradle

We have sketched out an argument that makes the behavior of the Newton's cradle plausible, in particular we addressed the question of why some situations that would conserve energy and momentum do not occur. But we haven't looked at what's going on within the string of balls to bring about this result.

1. The successive impacts model. The simplest model to understand is one that invokes a "cheat". It assumes the N balls are initially **not** touching. The first ball is pulled back and strikes the second with speed. The first ball comes to rest and the second moves forward with speed V , hits the third ball, and so on down the line, till the last ball is ejected with speed V .

This is valid when the balls are actually separated. But then some folks assume that the explanation is also valid when the balls are touching. Well, the results are nearly the same in both cases, but the dynamics of the processes are certainly different. We will move on to look at the interesting case, where all balls are initially touching each other.

2. The compression pulse model. This assumes that a compression pulse begins in the metal balls at the point of first impact, traveling through the balls with the speed of sound. The speed of sound in the material of which the balls are made is much greater than the speeds of the balls in this device. So the pulse "does its work" before any of the stationary balls have moved. The pulse travels forward and backward, reflecting from the ends of the string of balls and meeting again simultaneously at one point. Where is that point? Well, if the pulse originated between the first two balls, the pulse meets between the last two balls, where it gives up its energy to the last ball.

This sounds plausible at first. But there's a troublesome issue. This model requires that **all** of the energy of the first ball is given up to a compression pulse and **all** of that energy ends up and one localized point at the point where the last two balls touch. How does it **do** that without dispersal of energy, for the pulse initially goes in all directions, forward, backward, up, down and all directions in between? Pulse energy is reflected from the ball surfaces (the balls are round after all) in very complex paths (and most of these paths are not equal in length from start to finish). Though it sounds good, it fails to convince the skeptical student.

How does this pulse give up *all* of its energy to the last ball, instead of giving some back to the string of balls behind it?

Does this model hold up if the balls have identical mass but different sizes? What if a larger ball were placed in the middle of the array, say twice the diameter, replacing two of the other balls? This should not affect the time of arrival of the

pulse at the other end. But it certainly does affect the outcome of the experiment. Perhaps simpler, add mass to a ball in the middle of the array, by attaching a weight to its bottom. Why should this affect the compression pulse? But it does affect the outcome.

3. The balls-and-springs model. This model imagines a linear string of balls with small springs between them. It treats the system as a lattice array. It turns out that to make this work as a simulation of the spherical ball Newton's cradle, the springs will not obey Hooke's law, $F=-kx$, but rather obey the Hertzian spring law, $F = -kx^{3/2}$. This, it is argued, is a result of the balls being spherical. A linear array of objects of different shape, say cylinders, would behave differently. While interesting, this model is not an exact simulation, for its predictions do not quite match the real behavior.

4. The interface compression model. We know that at impact the balls, being elastic bodies, must deform at the point of impact. This stores compressional energy. In the N-ball case we assume that all balls are in contact and moving at nearly the same speed. So conservation of momentum allows us to calculate that speed. (In the case of N balls with one ball initially moving at speed V, the balls move at V/N during the contact phase.) Then we can calculate the kinetic energy of this composite body. It will be less than the initial kinetic energy of the system, because some energy is temporarily stored as elastic energy at the interfaces. This energy is distributed at all of the interfaces, but not equally (except in a few special cases). During decompression, work is done as this energy is released at the interfaces. (In the case of one ball initially moving at speed, it is slowed to a stop, as are all the others except the last ball in the string, which moves away with speed V.)

It all seems plausible, but doesn't answer the question of how the stored energy is distributed at the N-1 interfaces. That depends on the initial conditions. Consider three balls, with ball 1 moving at speed V, ball 2 at rest, and ball 3 moving at speed -V. From the symmetry, we expect that the compressional energy is equal at the two interfaces. But in the case with ball 1 moving at speed V, with ball 2 and 3 at rest, we might suppose that twice as much energy is stored at interface 12 as at interface 23. Why? During the compression phase the force at interface 12 acts on the two other balls (total mass 2), accelerating them and doing work on them. The force at interface 23 acts on ball 3 (mass 1), accelerating it by nearly the same amount. So the force at 12 must be twice as large as the force at 23. Since these forces act over the same distance, the work they do is also in ratio 2:1. The stored energies at the interfaces are therefore in 2:1 ratio.

But if the initial situation were speed V for ball 1, zero for ball 2 and -V for ball 3, then the forces at the interfaces are equal, and the energies stored there are equal.

Testing the models. Any model we might devise ought to be successful for the N-ball case with identical spherical elastic balls. It also ought to be successful in the

case where the balls are of mixed sizes, shapes, and masses. The very fact that new models are proposed every year, in professional journals, is evidence that there's no fully successful model yet, certainly no simple one suitable for elementary physics classes. The bibliography references at the end of this web page is further evidence of that.

Asymmetric Balls.

A particularly interesting variation of this toy is made with several equal balls, but including one ball of larger mass than the others. We called this the "oddball cradle". Consider the three-ball version:

1. **o oo**, One moving two stationary, equal mass (for comparison)
2. **O oo**, Large ball moving, two small ones stationary.
3. **o Oo**, Small one moving, other two stationary, large one in the center.
4. **o oO**, Small one moving, other two stationary, large one on other end.
5. **O oO**, Large one moving, other two stationary, large one on other end.

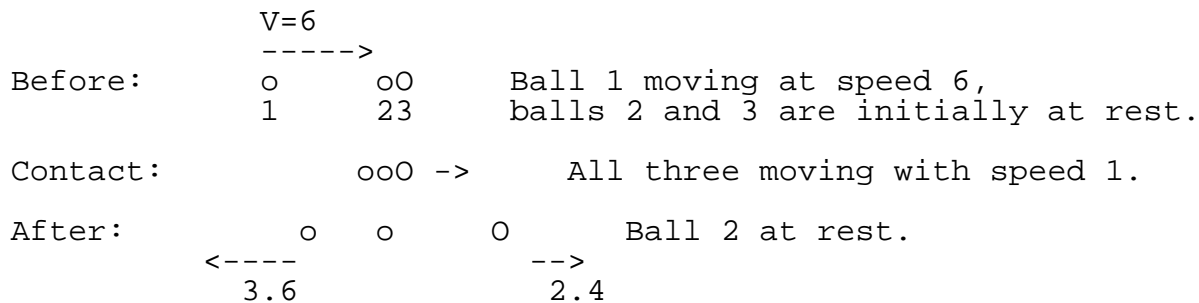
Can you predict the outcomes? Can you predict which ball moves faster, and whether one or more balls remain stationary after impact. Here's an [AVI movie of case 4](#). It's size is 600 kB. You need the Windows media player or some other AVI player to view this video clip.

Answers: Doing the experiment we observe that cases 1 and 5 have the same outcome: the middle ball remains at rest after collision and the last ball in line comes out with the same speed the first one had. It's as if the middle ball didn't participate. In case 2 all three balls move forward after collision. In case 3, the first ball rebounds, the other two move forward, the smallest one moving fastest. How well did you do?

A naive conclusion is that in any 3 ball collision with the middle ball initially at rest, the middle ball will remain stationary after collision, the other two behaving exactly as in the two ball case when they hit each other. In cases 1 and 5 that's exactly what happens. But in case 2 it does not. Case 2 seems to be the reverse of case 4, but its outcome isn't. There's a profound clue in these results.

Remember that we noted above that the 2-ball case is such that only one solution results from application of the laws of energy and momentum conservation. The predicted result is the one observed. Also, Newton's cradles with initial separation between the balls, produce nearly the same observed results as the kind with balls in contact, and may be modeled as successive 2-ball collisions. This gives important insight into the case where the balls are initially touching. If we model this case by saying that separation occurs first at the left-most interface, and then successively at the other interfaces, we can treat this problem as a succession of two-body interactions, and such a model predicts one unique outcome, which is the one observed.

Commercial versions sold as scientific apparatus for physics demonstrations often include one ball that is 3 or 4 times the mass of the others. What do you suppose this is intended to demonstrate? In case 4 above we can calculate, by the successive collisions model what the outcome should be:




This is the observed outcome, as seen in the video above and summarized below.

Case 4, masses 1, 1 and 4.

Quantity	Initial 1	Initial 2	Initial 3	Final 1	Final 2	Final 3	Net
Masses	1	1	4	1	1	4	6
velocity	6	0	0	-3.6	0	2.4	
momentum	6	0	0	-3.6	0	9.6	6
kinetic energy	18	0	0	6.48	0	11.52	18

Final exam: Let's look at the second case above. Check these results.



This is summarized below, calculated by the successive collisions model.

Masses 4, 1 and 1.

Quantity	Initial 1	Initial 2	Initial 3	Final 1	Final 2	Final 3	Net
Masses	1	1	3	1	1	3	6
velocity	6	0	0	2.16	5.76	9.6	

momentum	24	0	0	8.64	5.76	9.6	24
kinetic energy	72	0	0	9.3312	16.5888	46.08	72

Now do the experiment to see whether the actual device, with balls initially touching, matches these results from the successive collisions model. Any model of these devices must successfully predict situations with unequal mass balls as well as with unequal size and shapes.

Bibliography.

These are all from the journals:

American Physics Teacher (APT)
 American Journal of Physics (AJP)
 The Physics Teacher (TPT)

- Harvey B. Lemon, An Almost Forgotten Case of Elastic Impact, TPT 3, 36 (1935). [Discussion of the 3:1 mass ratio case.]
- Seville Chapman, A problem Involving the Order of Impacts, Proceedings of the AAPT Philadelphia meeting, December 1940, AJP 9, 56 (1941).
- Seville Chapman, Some Interesting Aspects of the Collision Ball Apparatus, AJP 9, 357-360 (1941).
- Harold K. Schilling and Henry Yeagley, A New Impact Apparatus, AJP 15, 60-64 (1947).
- John V. Kline, The Case of the Counting Balls, AJP 28, 102-103 (1960).
- Seville Chapman, Misconception Concerning the Dynamics of the Impact Ball Apparatus, AJP 28, 705-711 (1960).
- R. Edwin Worley, Impact Demonstration with Plastic Croquet Balls, AJP 30, 769-770 (1962).
- James D. Kerwin, Velocity, Momentum, and Energy Transmissions in Chain Collisions, AJP 40, 1152-1157 (1972).
- Leonard Flansburg and Karl Hudnut, Dynamic solutions for linear elastic collisions, AJP 47, 911-914 (Oct 1979).
- Alexander Langsdorf, Jr., Letter: Referees Beware!, AJP 48, 335 (1980).
- Christian Lehmann and Manfred Rosenbauer, Erratum: "Focusing and defocusing collisions of atoms in solids: A classroom demonstration," [AJP 48, 496-497 (1980)], AJP 49, 89 (1981)..
- F. Herrmann and P. Schmalzle, Simple explanation of a well-known collision experiment, AJP 49, 761-764 (1981).
- Alexander Langsdorf, Jr., Letter: AJP 50, 105 (1982).
- F. Herrman and M. Seitz, How does the ball-chain work?, AJP 50, 977-981 (1982).
- Jean C. Piquette and Mu-Shiang Wu, Comments on "Simple explanation of a well-known collision experiment," AJP 52, 83 (1984).
- F. Herrmann and P. Schmalzle, Response to "Comments on "Simple

- explanation of a well-known collision experiment," AJP 52, 84 (1984).
- Matthias Reinsch, Dispersion-free linear chains, AJP 62, 271-278 (1994).
 - David Auerbach, Colliding rods: Dynamics and relevance to colliding balls, AJP 62, 522-525 (1994).
 - Robert Erlich, Experiments with "Newton's Cradle," TPT 34, 181-183 (March 1996).
 - Roura, Pere, Collision duration in the elastic regime," TPT 35, 435-436 (Oct 1997).
 - Gavenda, J. D. and J. R. Edington, "Newton's Cradle and Scientific Explanation," TPT, 35, 411-417 (Oct 1997).
 - Mascarenhaus, F. M. F., C. M. Spillmann, J. F. Lindner, and D. T. Jacobs, "Hearing the shape of a rod by the sound of its collision," AJP, 66, 692-697 (Aug 1998).
 - Roura, Pere, Collisions Between Rods: A Visual Analysis," TPT 41, 32-35 (Jan 2003).
 - J. D. Gavenda and J. R. Edington, "Newton's cradle and scientific explanation," TPT 35, 411 (1997).
 - [Pira Bibliography, linear momentum and collisions.](#)
-

July, 2001. Revised 2002, 2003, 2003. Major revisions, August 2014)

Input and suggestions are welcome at the address to the right. When commenting on a specific document, please reference it by name or content.

DSIMANEK @ LHUP.EDU

Return to [Donald Simanek's page.](#)