

Original citation:

Sanborn, Adam N., Mansinghka, Vikash K. and Griffiths, Thomas L. (2013) Reconciling intuitive physics and Newtonian mechanics for colliding objects. *Psychological Review*, 120 (2). pp. 411-437.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/54487>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

© APA 2013

<http://dx.doi.org/10.1037/a0031912>

"This article may not exactly replicate the authoritative document published in the APA journal. It is not the copy of record."

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP URL' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Reconciling Intuitive Physics and Newtonian Mechanics for Colliding Objects

Adam N. Sanborn

University of Warwick

Vikash K. Mansinghka

Massachusetts Institute of Technology

Thomas L. Griffiths

University of California, Berkeley

Send Correspondence To:

Adam Sanborn

Department of Psychology

Coventry CV4 7AL, UK

a.n.sanborn@warwick.ac.uk

Supplemental code: http://www2.warwick.ac.uk/fac/sci/psych/people/academic/asanborn/asanborn/psych_review_code.zip

Author Note

This work was supported by an NSF Graduate Research Fellowship, a Royal Society USA Research Fellowship, and the Gatsby Charitable Foundation (ANS) and grant number FA9550-07-1-0351 from the Air Force Office of Scientific Research (TLG). Preliminary results were presented in the Proceedings of the 31st Annual Conference of the Cognitive Science Society. The authors thank Patricia Cheng, Michael Lee, and three anonymous reviewers for helpful discussions.

Abstract

People have strong intuitions about the influence objects exert upon one another when they collide. Because people's judgments appear to deviate from Newtonian mechanics, psychologists have suggested that people depend on a variety of task-specific heuristics. This leaves open the question of how these heuristics could be chosen, and how to integrate them into a unified model that can explain human judgments across a wide range of physical reasoning tasks. We propose an alternative framework, in which people's judgments are based on optimal statistical inference over a Newtonian physical model that incorporates sensory noise and intrinsic uncertainty about the physical properties of the objects being viewed. This "noisy Newton" framework can be applied to a multitude of judgments, with people's answers determined by the uncertainty they have for physical variables and the constraints of Newtonian mechanics. We investigate a range of effects in mass judgments that have previously been taken as strong evidence for heuristic use and show that they are well explained by the interplay between Newtonian constraints and sensory uncertainty. We also consider an extended model that handles causality judgments, and obtain good quantitative agreement with human judgments across tasks that involve different judgment types with a single consistent set of parameters.

Reconciling Intuitive Physics and Newtonian Mechanics for Colliding Objects

People believe that they understand how everyday physical objects behave. However, our intuitive understanding appears to be inconsistent with Newtonian mechanics: People often predict that an object that is swung around will follow a curved path when released (McCloskey, Caramazza, & Green, 1980) and predict that a ball dropped from a moving object will fall straight downwards (McCloskey, Washburn, & Felch, 1983; Kaiser, Proffitt, Whelan, & Hecht, 1992). Following the ground-breaking work of Michotte (1963), collisions between objects have been used as the basis for one of the most comprehensive investigations of the relationship between intuitive and Newtonian mechanics (A. Cohen, 2006; A. Cohen & Ross, 2009; Gildea & Proffitt, 1989, 1994; Runeson, 1983, 1995; Runeson & Vedeler, 1993; Runeson, Juslin, & Olsson, 2000; Schlottmann & Anderson, 1993; Todd & Warren, 1982). In a typical experiment, participants observe a collision between two objects and are then asked to make a judgment about the physical properties of the objects involved (such as their relative mass) or the relationships between them (such as whether one object caused the other to move). Analysis of these judgments has revealed significant deviations from the predictions of Newtonian mechanics, which have been held up as evidence that intuitive physics is based on a set of shortcuts or heuristics (Andersson & Runeson, 2008; A. Cohen, 2006; A. Cohen & Ross, 2009; Gildea, 1991; Gildea & Proffitt, 1989, 1994; Michotte, 1963; Runeson et al., 2000; Schlottmann & Anderson, 1993; Todd & Warren, 1982). In this paper, we show that dissociations between intuitive physics and Newtonian mechanics can be reconciled using a rational model that takes into account uncertainty about sensory information.

Judgments of the relative masses of objects involved in collisions have been taken as providing some of the strongest evidence for heuristic accounts. In the standard mass ratio judgment task two objects collide with each other on a screen and observers must judge which object was heavier (see Figure 1A). Newtonian physics provides a simple solution to this problem: If the initial and final velocities of the objects are known, then the object with the smaller difference between initial and final velocity has the greater mass. However, people show sensitivity to irrelevant information, such as whether kinetic energy is conserved, as well as

systematic biases, such as a strong bias toward believing that the object with higher initial speed is heavier (Andersson & Runeson, 2008; A. Cohen & Ross, 2009; Gilden & Proffitt, 1989; Runeson & Vedeler, 1993; Runeson et al., 2000; Todd & Warren, 1982). These deviations from a Newtonian account have been explained in terms of heuristics: looking for the object with the higher final speed and looking for the presence of a ricochet of one object off another, both of which are indications that this object is lighter (Gilden & Proffitt, 1989; Proffitt & Gilden, 1989; Todd & Warren, 1982).

Proceeding almost independently (but see Kaiser & Proffitt, 1984), research in collision detection has also produced heuristic models of the formation of causal impressions. In this task, one object moves toward a stationary second object, and participants are to say whether a collision occurred or whether the second object moved on its own. Researchers have manipulated variables such as the time delay between movements, the gap between objects, and the relative velocities to understand how people infer causality (e.g., Boyle, 1960; Choi & Scholl, 2004; L. B. Cohen & Oakes, 1993; Gemelli & Cappellini, 1958; Guski & Troje, 2003; Leslie, 1984; Leslie & Keeble, 1987; Michotte, 1963; Natsoulas, 1961; Parovel & Casco, 2006; Schlottmann & Anderson, 1993; Schlottmann & Shanks, 1992; Straube & Chatterjee, 2010). Larger gaps and longer time delays reduce the impression of causality as would be expected if the impression of causality arose from an internal model of Newtonian mechanics. However, Michotte argued that people also show discrepancies between intuitive and Newtonian mechanics (Michotte, 1963, p. 87, 220), finding a dissociation between causal impression and the physical effect the moving object had on the stationary object. The dissociation came from a comparison of two cases. In the first case, the initially stationary object was faster after contact than the initially moving object was before contact. In the second case, the initially stationary object was slower after contact than the initially moving object was before contact. It is in the first case that the effect of the initially moving object was greater, but participants reported a greater causal impression in the second case (Michotte, 1963; Natsoulas, 1961). As a result, Michotte argued that people do not learn to judge causality from experience.

The arguments made against a Newtonian account of judgments of relative mass have generally assumed that at some level the visual system accurately detects the velocities of the objects involved in a collision, which provide information about their relative mass. This is consistent with Gibson's (1966) doctrine of "direct perception," which plays a central role in the paradigm of ecological psychology in which many of these experiments were conducted (Runeson et al., 2000; Todd & Warren, 1982). However, recent work in visual perception has emphasized the inductive challenge posed by noisy sensory stimulation, and shown that many puzzling aspects of human vision can be explained as rational Bayesian inference using a probabilistic model calibrated to the visual environment (Geisler, Perry, Super, & Gallogly, 2001; Kersten, Mamassian, & Yuille, 2004; Najemnik & Geisler, 2005; Schwartz, Sejnowski, & Dayan, 2009; Weiss, Simoncelli, & Adelson, 2002; Yuille & Kersten, 2006).

In this paper, we apply this probabilistic approach to intuitive physics, showing that people's judgments about the relative masses of objects are consistent with Newtonian mechanics once uncertainty about the velocities of the objects is taken into account. The resulting rational model of intuitive dynamics can also be used to model decisions about whether one object caused another to move (see Figure 1B), providing the first theoretical account that unifies recent work on judgments of relative mass with Michotte's (1963) classic research on perceptual causality. Finally, we discuss how Newtonian mechanics could be learned, how this framework can explain a host of effects on how judgments of one variable depend on the others, what sort of intuitive knowledge about mechanics people seem to have, and the future challenges for this framework.

Collisions and Heuristics

How people perceive colliding objects has driven a long line of inquiries into the workings of the mind. When Hume (1748) explained why we could not directly reason from causes to effects, he used billiard balls as an example. When a moving billiard ball collides with a stationary billiard ball, there is a wealth of possible outcomes: the moving ball could stop dead, it could bounce backwards, or could carry off in any direction. Hume argued that without

experience we would not be able to know which outcome would occur and that our experiences were “conjoined but never connected” (section VII, part ii of Hume, 1748).

Our experiences of colliding objects in the world are well described by Newtonian mechanics, which constrain the possible outcomes for the two billiard balls. When two billiard balls (or cubes, as shown in Figure 1A) without any spin have a head-on impact, we can determine their final velocities by treating the problem as the collision of two point masses. The objects a and b with masses m_a and m_b move with initial velocities u_a and u_b . They participate in a collision and then move apart with final velocities v_a and v_b . The final velocities are given by

$$v_a = \frac{m_a u_a + m_b (u_b + e(u_b - u_a))}{m_a + m_b} \quad (1)$$

$$v_b = \frac{m_b u_b + m_a (u_a + e(u_a - u_b))}{m_a + m_b} \quad (2)$$

where e is the coefficient of restitution for the collision, with $e = 1$ indicating that kinetic energy is conserved, approximated by a collision between billiard balls, while $e < 1$ indicates that some kinetic energy has been converted into heat or object deformation, as in a collision of clay balls.

Michotte (1963) believed that Hume had asserted that we could not perceive a tight causal connection between events and took issue with this claim. In a series of over a hundred experiments, Michotte experimented with colliding objects that closely matched Hume’s example, in which one object, termed the *motor object*, was initially in motion and the other, the *projectile object*, was initially stationary. He demonstrated that with many realistic collision events, participants would have a strong and immediate impression of causality when one object collided with another. For these realistic collision events, such as those determined by the above equations, participants tended to report a causal impression of *launching*: The motor object had launched the projectile object.

Michotte also experimented with displays that were not realistic collisions, introducing in some cases a gap, g , between objects so that they never touched, or in others a time delay, t , so that there was a pause between the cessation of the first object’s motion and the beginning of the second object’s motion. A schematic of the variables in this type of trial is shown in Figure 1B.

Michotte found many close correspondences between his participants' impressions of causality and the results of Newtonian mechanics. Introducing a long time delay between object *a* stopping and object *b* starting or a large gap between the two objects destroyed the impression of causality. The velocity of the objects was also important. When the projectile object moved off at an implausible speed, causal impressions were greatly reduced.

Despite the close correspondence between causal impressions and Newtonian mechanics, Michotte did not believe that causal impressions were learned from experience with colliding objects. His rejection of learning was due to his finding “innumerable” discrepancies between the displays that participants received a causal impression from and Newtonian mechanics. One example he featured was the dependence of causal impression on the relative velocity of the motor and projectile objects. In his Experiment 39, in which the projectile object traveled slower than the incoming motor object, the causal impression of launching was strong. But for his Experiment 40 in which the projectile object traveled faster than the incoming motor object, the causal impression of launching was weak. There was a dissociation between the strength of the causal impression and the strength of the effect of the motor object on the projectile object. Michotte considered this a paradox: How could the impression of causality be greater when the effect on the projectile object was less?

Michotte's work inspired a large number of studies investigating the causal impressions generated by displays of moving objects. Many later researchers disagreed with his claim that people were not learning their intuitive mechanical expectations from the environment (Dittrich & Lea, 1994; Rips, 2011; Weir, 1978; White & Milne, 1997), including Runeson (1983) who criticized the whole approach of investigating people's causal impressions as misguided. He dismissed causality judgments as subjective, and following Gibson's ecological psychology instead emphasized properties of the environment that people would have a use for, such as the mass of objects (Gibson, 1966).

Runeson (1983) named his approach the Kinematic Specification of Dynamics (KSD). The variables involved in a collision can be divided into the *kinematic* observable variables, such as

velocities, and the *dynamic* variables, such as mass, that determined the object paths. The kinematic variables limited the possible values of the dynamic variables of a display, but the dynamic variables could not always be determined uniquely. Runeson showed that while the specific masses of objects were not determinable from observing the kinematic behavior of objects, the ratio of the masses was. Using the conservation of momentum, he derived

$$\frac{m_a}{m_b} = \frac{u_b - v_b}{v_a - u_a}. \quad (3)$$

This relationship gives a simple “invariant” of the kind that Gibson (1966) suggested might support direct perception of properties of the world.

Runeson realized the KSD made strong predictions for people’s behavior. He proposed that the account be tested by having participants judge the mass ratio of objects from physically realistic displays while the coefficient of restitution was manipulated. Equation 3 does not involve e , and thus judgments should be invariant to changes in the coefficient of restitution. This prediction was first tested by Todd and Warren (1982), using a design in which both objects were in motion before contact. They found that the coefficient of restitution had a significant effect on the accuracy of participants’ mass judgments, with lower values of e leading to lower accuracy in identifying which object was heavier.

The second experiment of Todd and Warren (1982) revealed an even more puzzling and striking apparent departure from Newtonian mechanics. The stimuli in this experiment were similar to Michotte’s canonical launching display, in which an initially moving motor object collided with an initially stationary projectile object. If the mass of the motor object was greater, participants responded accurately. However, if the mass of the projectile object was greater, participants often still reported that the motor object was heavier. Participants’ point of subjective equality of the masses of the two objects was shifted: There was a bias towards believing that the initially moving object was heavier, which came to be known as the *motor object bias*.

The motor object bias has been replicated and extended in many other studies (Andersson & Runeson, 2008; A. Cohen, 2006; A. Cohen & Ross, 2009; Flynn, 1994; Gilden & Proffitt, 1989; Jacobs, Michaels, & Runeson, 2000; Jacobs, Runeson, & Michaels, 2001; Runeson et al.,

2000; Runeson & Vedeler, 1993). It has been taken as strong evidence against KSD and more generally as evidence that people cannot accurately judge the relative mass of colliding objects (Gilden, 1991). In the place of accurate perception of mass, it has been proposed that collisions between two objects present too complicated a problem for people to analyze correctly. Instead people can only use restricted heuristics for judging mass, such as which object moves faster after a collision or which object ricochets (A. Cohen, 2006; A. Cohen & Ross, 2009; Gilden & Proffitt, 1989; Proffitt & Gilden, 1989; Runeson et al., 2000).

Combining Noisy Perception, Probabilistic Inference, and Newtonian Mechanics

Marr (1982) was sympathetic to Gibson's claim that people used perceptual invariants, such as the mass ratio in the KSD approach, but believed Gibson had greatly underestimated the difficulty required to construct invariants. Despite our apparent ease in interpreting the environment, the mapping between sensory information and interpretation is extremely complex. Marr emphasized this difficulty and put forward the idea of analyzing cognitive systems by focusing on the computational problems posed by the environment: What problem is the cognitive system trying to solve?

Here we focus on the problem of drawing inferences based on noisy sensory data, as even our guesses about relatively simple perceptual variables, such as velocity, distance, and time, are not perfect (Andrews & Miller, 1977; Ekman, 1959; Hick, 1950; Notterman & Page, 1957). The problem of finding the best possible interpretation of noisy sensory information is one that we can solve by following the rational analyses of Shepard (1987) and Anderson (1990) and the ideal observer in visual and auditory research (de Vries, 1943; Geisler, 1989; Green & Swets, 1966; Peterson, Birdsall, & Fox, 1954; Rose, 1942). The best possible interpretation depends on the noisy information observed O , the observer's beliefs about what states S are likely to occur in the environment, $p(S)$, and knowledge of the noise process: How frequently each possible observation arises from each individual state of the environment, $p(O|S)$. Given this information, we can find how likely each state of the environment is given the observations we have by using

Bayes rule,

$$p(S|O) = \frac{p(O|S)p(S)}{\sum_{S'} p(O|S')p(S')}. \quad (4)$$

where the denominator sums over all possible states of the environment.

This approach has been used successfully in a wide range of settings (Geisler et al., 2001; Kersten et al., 2004; Najemnik & Geisler, 2005; Schwartz et al., 2009; Weiss et al., 2002; Yuille & Kersten, 2006) to explain how people make inferences about what object has generated visual sensory information. Here we propose that this approach can be extended to predicting people’s judgments about physics. Like in previous work (McIntyre, Zago, Berthoz, & Lacquaniti, 2001; Shepard, 1984; Zago, McIntyre, Senot, & Lacquaniti, 2009), we assume that people have an internal representation of physical constraints. We then assume that people appropriately combine this prior knowledge that is constrained by Newtonian mechanics with noisy information from the sensory system to make inferences about the physical situation. This idea was foreshadowed by Watanabe and Shimojo (2001) who proposed that people combining visual and auditory information to form a causal impression were solving the problem of “inverse physics” – attempting to work out not just what object is generating the sensory information, but the physics governing the objects as well. This approach, which here we call *noisy Newton*, was introduced in Sanborn, Mansinghka, and Griffiths (2009) and has been subsequently applied to other aspects of intuitive physics (Hamrick, Battaglia, & Tenenbaum, 2011; Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2012; Smith & Vul, 2012)

The noisy Newton framework provides a model of how observed variables are generated based on states of the world, defining $p(O|S)$ in Equation 4. This model can be used to work backward from the observed variables to answer questions about the state of the world by using Bayes’ rule to calculate $p(S|O)$. Specific queries can be addressed by summing the posterior probabilities of the states that are consistent with that query. For example, we might have observed noisy versions of the velocities of objects in a collision (our observed variables), infer their masses and the coefficient of restitution (our posterior probability distribution), and then use

the resulting distributions on the masses to evaluate whether one object weighs more than the other by summing the probabilities of the states where this is true (answering our specific query). This combination of Newtonian physics and Bayesian inference can potentially be applied to any problem that requires making an inference from observed variables.

The power of the noisy Newton framework is that it can unify many different types of intuitive physical judgments. In this paper, we use it to examine people's intuitions about the mechanics of two colliding objects in a one-dimensional collision. This case has a relatively simple mechanical description, but people's intuitions about these stimuli have generated a lot of interest among psychologists. In the following sections, we introduce extant models of mass judgments and explore how well noisy Newton compares to these models in predicting the relationship between behavior and Newtonian mechanics. Next, we demonstrate the power of the method by extending it to predict how the impression of causality depends on the physical variables. Next, we demonstrate how the noisy Newton framework does not necessarily require explicit or innate knowledge of physical constraints, but could be learned from a reasonably small number of examples. Finally, in the General Discussion, we explore the larger body of research on the relationship of how estimates of physical variables depend upon the display, explore what sort of mechanical knowledge we might intuitively have, and examine the challenges going forward for this approach.

Judgments of Mass

Evaluating the relative mass of objects is a well-studied intuitive physics task that has provided strong evidence in favor of heuristic models. We begin by outlining two models that have been used before for mass judgments as well as the form of the noisy Newton model for mass judgments. Then we apply these three models to the effects of the coefficient of restitution, the motor-object bias, the effect of occlusion, and the effect of training.

Models of Mass Ratio Judgments

Direct Perception. The *direct perception* model follows from Runeson (1983)'s KSD approach. It assumes that people have direct access to the initial and final velocities, allowing them to compute the mass ratio in Equation 3 from noiseless perception of the displayed velocities. As a result, identification of the object with greater mass is invariant to the coefficient of restitution. This is the model that has been used to evaluate whether people make judgments that are consistent with Newtonian physics in the past (A. Cohen & Ross, 2009; Gilden, 1991; Gilden & Proffitt, 1989, 1994; Runeson, 1983, 1995; Runeson & Vedeler, 1993; Runeson et al., 2000; Todd & Warren, 1982), although we use a more flexible version that allows noise following the computation of the mass ratio to provide a closer comparison to our other models (A. Cohen, 2006; Runeson et al., 2000).

Heuristic. The *heuristic* model was developed under the assumption that people's intuitions correspond with Newtonian mechanics in very simple displays, but the relationship governing two colliding objects is too complex to understand accurately. Instead, people used heuristics: simpler quantities that approximate the mass ratio, but not perfectly. It was hypothesized that people combine two heuristics to make a decision: the object with greater final speed is lighter, and an object that ricochets is lighter than one that does not (Gilden & Proffitt, 1989; Gilden, 1991; Proffitt & Gilden, 1989).

These heuristics were proposed as a qualitative alternative to a Newtonian account, and the salience functions that determine when the heuristics are used was left undefined. We follow A. Cohen (2006), and generate quantitative predictions from this approach by fitting salience thresholds to the data. If only one of the heuristics is salient, then the response indicated by the salient heuristic is the response made by the model. If both heuristics are salient, then one of the heuristics, the default, overrules the other. If neither heuristic is salient, then the model randomly selects a response with equal probability (Gilden & Proffitt, 1989). We allow the default heuristic to be selected probabilistically to approximate individual differences and use a single salience threshold for all participants, though we consider more flexible formulations below.

Noisy Newton. Finally, the *noisy Newton* model treats the masses of the objects and the coefficient of restitution as hidden variables that need to be inferred from noisy observations of the initial and final velocities. Combining the observed velocities with prior beliefs, we can use Bayesian inference to calculate the probability that m_a is greater than m_b , making it possible to select the object with greater mass. A graphical model that explains the dependencies between the variables is shown in Figure 2.

We set simple prior belief distributions on the masses, coefficient of restitution, and the initial velocities. The initial velocities were assumed to follow a normal distribution with a mean of zero and a standard deviation set to match human responses. The masses were assumed to follow an exponential distribution, so that lower masses were considered more likely than higher masses. Because the final velocities only depend on the ratio of the masses, the parameter of the exponential distribution has no effect on the calculation and the prior does not depend on the units used for mass. The coefficient of restitution was given a uniform prior distribution over its entire range from zero to one. Finally, the prior distributions on the final velocities followed Newtonian mechanics: they were calculated from the prior distributions on the other variables using Equations 1 and 2.

The noise in each variable was psychophysically motivated. The observed values were the true underlying Newtonian mechanical values combined with noise that approximately followed Weber's law: a standard deviation of the noise increasing linearly with the value of the variable. We set two parameters of this noise to match human judgments. More details about the noisy Newton model are given in Appendix A.

Effect of Coefficient of Restitution

Runeson (1983) proposed an experiment to test the KSD approach: Determine whether people's mass judgments were invariant to changes in the coefficient of restitution. We replicated the test of this in Experiment 1 of Todd and Warren (1982) (details in Appendix B) and found the same result: A smaller coefficient of restitution resulted in less accurate judgments, as shown in

Figure 3. The coefficient of restitution, e , which was manipulated independently of the mass ratio, changed the accuracy of people's judgments, with better accuracy occurring with greater coefficients of restitution. A repeated-measures ANOVA showed a main effect of mass ratio, $F(3, 57) = 133, p < 0.001$, and a main effect of coefficient of restitution, $F(2, 38) = 16.0, p < 0.001$. There was a marginally significant interaction between mass ratio and coefficient of restitution, $F(6, 114) = 1.89, p = 0.08$.

The direct perception model, because it is invariant to changes of the coefficient of restitution, did not predict the separation between coefficient of restitution conditions. Models were evaluated by computing the negative log likelihood (NLL) of the human data from this experiment given the best-fitting model parameters. The best fitting parameters were calculated by fitting the model to the individual responses from this experiment and the replication of Experiment 2 of Todd and Warren (1982) that we present in the next section. We also computed Akaike Information Criterion (AIC; Akaike, 1973) values which combine the log likelihood measure of goodness of fit with a penalty for complexity based on the number of parameters in the model. For both measures, smaller values are better and a measure of the relative posterior probability of the models assuming equal prior probabilities can be computed from a transformation of their AIC values (Akaike, 1978; Burnham & Anderson, 2002; Wagenmakers & Farrell, 2004). For the direct perception model, $NLL = 999$ and $AIC = 1,999$.

The heuristic model was developed to account for this result, and relying on the final speed to determine which object is lighter predicts the difference in conditions ($NLL = 716$ and $AIC = 1,437$). The model produces a much better match to the human data than the direct perception model, and the separation in the predicted accuracy for the coefficient of restitution conditions matched human judgments both qualitatively and quantitatively.

The noisy Newton model also accounts for this difference ($NLL = 696$ and $AIC = 1,398$) and with a better quantitative fit than the heuristic model, despite using parameters that chosen to produce qualitative results and not the best quantitative fit: The relative probability of this model, compared to the direct perception and heuristic models using Akaike weights, is nearly one. We

can understand why the model makes this prediction by first observing that a decision boundary can be constructed in the space of the final velocities of the two objects by rearranging Equation 3 to obtain

$$v_a + v_b = u_a + u_b \quad (5)$$

as the set of velocities for which the masses of the two objects are equal. This equation forms a decision boundary between displays in which object *a* is heavier and displays in which object *b* is heavier; if the right side is larger than one object is heavier, but if the left side is larger than the other object is heavier. Figure 4A shows both the boundary and the velocities from the different coefficient of restitution conditions. From this figure, we can see that the conditions with low coefficients of restitution are closer to the decision boundary than conditions with a high coefficient of restitution. The same noise added to velocities closer to the decision boundary will result in lower performance than when it is added to velocities far from the decision boundary, and this is what drives the predictions of the noisy Newton model.

The robustness of the noisy Newton model to changes in its parameters was investigated by changing each parameter sequentially and investigating the effect on the predictions. We manipulated each of the three parameters above, the standard deviation of the prior on the initial velocities and the two noise perceptual noise parameters, to be one-fourth, one-half, double, and quadruple its initial value. In addition, because objects with very high coefficients of restitution are very rare, we manipulated the maximum of the uniform prior on the coefficient of restitution, so that it could range from zero to either 0.2, 0.4, 0.6, or 0.8 in addition to its default range of zero to one. Varying these parameters yielded sixteen additional sets of model predictions. Every parameter set showed the same ordering of accuracy of mass judgments by coefficient of restitution as the human data.

Motor Object Bias: Faster Moving Object Appears Heavier

While an effect of the coefficient of restitution was evidence against KSD, the second experiment of Todd and Warren (1982) found a more profound discrepancy between Newtonian

mechanics and mass judgments. People have a strong bias towards believing that the object that is initially moving faster is heavier, called the motor object bias. We replicated this experiment as well (details in Appendix B) and found the same result: People judged the initially moving object to be heavier when it had a higher mass, but also tended to judge it as heavier when it had a lower mass. Also, the motor-object bias interacts with the coefficient of restitution, producing a larger motor-object bias when the coefficient of restitution is lower.

These results are shown in Figure 5, plotting the proportion choice of object a against the mass ratio m_a/m_b . There was a clear bias towards believing $m_a > m_b$ even when the opposite was true. We tested whether there was a bias in response by fitting a cumulative Gaussian distribution with a constant probability of guessing to each coefficient of restitution condition. The best fitting points of subjective equality, the mass ratio at which the psychometric function crosses the indifference line of 0.5, were 0.91, 0.89, 0.59, and 0.27 for the coefficients of restitution $e = 1$, $e = 0.9$, $e = 0.5$, and $e = 0.1$ respectively. Confidence intervals for these points of subjective equality were bootstrapped by redrawing participants with replacement, fitting the psychometric function to each draw, and finding the smallest interval that covered the middle 95% of fitted values (Efron & Tibshirani, 1993). The upper bounds were 1.01, 1.00, 0.65, and 0.76 for the coefficients of restitution $e = 1$, $e = 0.9$, $e = 0.5$, and $e = 0.1$ respectively, demonstrating that the $e = 0.5$ and $e = 0.1$ conditions were reliably biased away from $m_a/m_b = 1$. The pattern of the results shows the point of subjective equality moving away from the objective point of equality as e decreases.

All three models used the same parameters for both mass judgment experiments, with the direct perception and heuristic models using the parameters that produced the best simultaneous fit to both experiments. Like the effect of the coefficient of restitution, the motor object bias cannot be explained by the direct perception model (NLL = 1,462 and AIC = 2,927). The mass ratio is correctly perceived with zero-median noise added, so the direct perception model's point of subjective equality is equal to the point of objective equality.

The heuristic account predicted some aspects of the human data well and other aspects not

so well, but was certainly quantitatively better than the direct perception model (NLL = 1,258 and AIC = 2,522). The ordering of the points of subjective equality for the heuristic model matched the empirical ordering, but otherwise the model was a poor match to the human data. A clear mismatch is the prediction for the $e = 0.1$ condition. We found the slope of a regression line between log mass ratio and proportion of choosing $m_a > m_b$ to be 0.17, and computed a confidence interval as above. The confidence interval of 0.10 to 0.23 did not include zero, indicating that the slope of the human data is positive, with the probability of choosing object a growing as the ratio m_a/m_b increases, as shown in Figure 2 of the main text. As described in Appendix C, even the slope predicted by the more general heuristic model with salience functions instead of thresholds cannot exceed zero for the range of mass ratios included in this experiment.

The noisy Newton model produced the best quantitative fit (NLL = 1,003 and AIC = 2,011): The relative probability of this model, compared to the direct perception and heuristic models using Akaike weights, is nearly one. Surprisingly it also produced the bias with the same parameters used to predict results for the replication of Experiment 1 of Todd and Warren (1982). The reason it did so was complex, though the overall intuition can be seen in Figure 4B. The prior distribution over initial velocities has reduced the inferred sum of initial velocities, and this shifts the decision boundary that determines the response of which object is heavier. The shift goes in a particular direction, so that most samples from a noisy distribution centered on the true values would fall on the wrong side of the boundary, leading to the motor object bias.

To produce the effect in this way, the perception of the sum of the initial velocities must shift more than the perception of the sum of the final velocities. This differential shift is due to the structure of the prior distribution, shown in Figure 6. The Newtonian calculation of $v_a + v_b$ from the random variables of initial velocities, masses, and coefficients of restitution resulted in a marginal prior distribution for $v_a + v_b$ with greater variance than the marginal prior distribution for $u_a + u_b$. We can see the effects of the prior distribution by calculating the posterior distributions for a trial in which $m_a = m_b$, $e = 1$, and the participant happens to sample the values $u_a = 3.18$ cm/s, $u_b = 0$ cm/s, $v_a = 0$ cm/s, and $v_b = 3.18$ cm/s, from a displayed collision. Here

the likelihood distributions for $u_a + u_b$ and $v_a + v_b$ are equal. If the prior distributions were equal for the initial and final velocities, then the posterior distributions would be equal as well, and the model would make the unbiased prediction of $m_a = m_b$. However, the influences of the prior velocity distributions that are peaked at zero are unequal: the narrower prior on $u_a + u_b$ causes the posterior distribution of $u_a + u_b$ to be pulled closer to zero than the posterior distribution of $v_a + v_b$. Figure 4B then shows how the greater shrinkage for $u_a + u_b$ results in the motor object bias.

As shown above, the key predictions are due to qualitative aspects of the model: worse performance for lower coefficients of restitution is due to noisy perception and the motor object bias is due to having a prior distribution over the sum of velocities with the mode at zero. We investigated the robustness of the noisy Newton model using the same sixteen additional simulations as described in the robustness analysis above. The subjective point of equality was found for each simulation: for 94% of the simulations the average subjective point of equality was biased and for 65% of simulations the subjective point of equality for every coefficient of restitution condition was biased. In contrast, the heuristic model could be augmented to have far more parameters than the noisy Newton model, but still miss qualitative aspects of the data in Experiment 2. In its original form (Gilden & Proffitt, 1989), each participant has their own salience function and default choice, but as demonstrated in Appendix C this would not allow it to fit the positive slope of the $e = 0.1$ condition in Experiment 2.

Motor Object Bias with Both Objects Moving

A possible explanation of the motor-object bias is that there is something special about a moving object colliding with a stationary object. Perhaps people are implicitly assuming that the motor object has to overcome the static friction of the stationary projectile object. However, the motor object bias also occurs when both objects are moving prior to contact (A. Cohen & Ross, 2009; Runeson et al., 2000; Runeson & Vedeler, 1993). Though not reported in Experiment 1 of Todd and Warren (1982), we found this bias in our replication when the data are plotted along the

axis of the ratio of the mass of the initially faster object over the mass of initially slower object as shown in Figure 7. Here the point of subjective equality between the two masses is shifted toward values where the slower object is heavier than the faster object, meaning that there is a bias towards perceiving the faster object as heavier. As before, the bias is greater for lesser values of e . We tested whether there was a bias in responses by fitting a cumulative Gaussian distribution with a constant probability of guessing to each coefficient of restitution condition. The best fitting points of subjective equality, the mass ratio at which the psychometric function crosses the indifference line, were 0.83, 0.81, and 0.78 for the coefficients of restitution $e = 0.9$, $e = 0.5$, and $e = 0.1$ respectively. The upper bounds of bootstrapped confidence intervals were 0.89, 0.91, and 0.90 for the coefficients of restitution $e = 0.9$, $e = 0.5$, and $e = 0.1$ respectively, showing that all conditions were significantly biased away from objective equality. The pattern of results show that the point of subjective equality moves away from the point of objective equality as the coefficient of restitution decreases.

The direct perception model again predicts no such bias, with the response proportion produced by the model depending purely on the ratio of the masses. The heuristic model predicts the same qualitative bias and provides a good quantitative fit. The noisy Newton model predicts this bias as well, for the same reasons it does when one object is initially stationary¹. The quantitative fit of the models is the same as reported above for the analysis of accuracy of Experiment 1 because these fits are for trial by trial data. A robustness analysis of the noisy Newton model generalizes these results. The subjective point of equality was found for each simulation: for all of the simulations the average subjective point of equality was biased and for 71% of simulations the subjective point of equality for every coefficient of restitution condition was biased.

¹The symmetry of the prior in Figure 6 causes the motor-object bias for whatever object is faster. Assuming u_b is negative and u_a is positive, object b is faster when $u_a + u_b < 0$. In this case, $u_a + u_b$ is pulled by the prior distribution in a more positive direction, causing the sum $u_a + u_b$ to be perceived as larger than $v_a + v_b$. This causes a bias towards believing the initially faster object b is relatively heavier than it is.

Effect of Occlusion

Another possible explanation for the motor-object bias is that people are somehow drawn to whatever object is initially faster and select it as the heavier object. A straightforward prediction of this explanation is that if the initial velocities of the two objects are occluded, then the motor object bias should be eliminated. In opposition to this prediction, the bias was found to be stronger when the initial velocities were occluded (Runeson & Vedeler, 1993). There is disagreement about whether the heuristic model predicts this result (Gilden & Proffitt, 1994; Runeson, 1995), but the noisy Newton model does predict it. Figure 8A shows the noisy Newton model predictions for Experiments 1 and 2 of Todd and Warren (1982) if only the final velocities are observed, which shows a stronger bias than if both initial and final velocities are observed (shown in Figures 3 and 5). These predictions arise because the inferred initial velocities are no longer constrained by the observed values, allowing the inferred values to be closer to mode of the prior belief at zero. The smaller inferred sum of $u_a + u_b$ relative to the inferred sum of $v_a + v_b$, which is still constrained by observed values, causes a larger bias in the predicted choices. The examination of the robustness of the noisy Newton model supports this interpretation strongly: for 94% of the parameter sets in Experiment 1 and all of the parameter sets in Experiment 2, the occluded task resulted in a larger motor object bias for every coefficient of restitution condition.

Occlusion studies highlight an advantage the noisy Newton model has over the direct perception and the heuristic models. Both the direct perception and the heuristic models cannot be naturally applied to displays in which any of the variables that they require to compute a response are occluded. In the case of direct perception, the initial velocities are necessary, so occluding them means it is unclear what the model would predict. In contrast, the noisy Newton model is robust to missing information: If aspects of the display are uninformative, prior information is used in their place. Further experimental studies of missing information may provide a way to test these different types of models.

Effect of Training

Previous work has shown that the motor object bias is greatly reduced by training (Jacobs et al., 2000; Runeson et al., 2000). A single set of heuristics does not predict this result. Instead, this reduction in bias with training has been taken as evidence that inaccurate heuristics are initially used, but with experience the correct combination of cues is eventually found. While such a transition describes the data, how heuristics are generated and selected is unclear.

Instead of hypothesizing qualitative shifts between types of combination rules, we can naturally integrate training trials into the prior distribution of the noisy Newton model. As a test, we can see what predictions the noisy Newton model would make if training had overwhelmed the pre-experimental prior distribution, so that the trained prior distribution matches the set of training trials. In studies that have shown the reduction of the motor object bias with training, two distinct and easy-to-discriminate sets of initial velocities relevant for the mass ratio calculation were used in both training and test trials (Jacobs et al., 2000; Runeson et al., 2000). We can simulate the effect that training with non-confusable initial velocities would have on noisy Newton model for Experiments 1 and 2 of Todd and Warren (1982). We ensure that during the test trials only the correct initial velocities from the prior distribution are used. If this is done, we find a reduced motor object bias predicted for both experiments, as shown in Figure 8B, because inference about the sum of initial velocities, $u_a + u_b$, which determine the decision boundary, is unbiased by the prior. The examination of the robustness of the noisy Newton model supports this interpretation: for 82% of the parameter sets in Experiment 1 and 65% of the parameter sets in Experiment 2, the training task resulted in a smaller motor object bias for every coefficient of restitution condition.

Impressions of Causality

As a further means of evaluating the framework, we considered how it could apply to the problem of identifying causality in moving objects, specifically one of the effects that Michotte identified as a dissociation between causality and Newtonian mechanics. While research on

judgments of relative mass was inspired by Michotte's (1963) experiments on causality, existing models of mass judgment do not make predictions about how people should decide whether one object caused another object to move. The strength of our framework is that we can develop models that are consistent between tasks. In this section, we develop the noisy Newton model for causality judgment tasks. We then apply it to a dissociation noticed by Michotte: When the projectile object is hit, the impression of causality is stronger if it moves more slowly than the motor object than if it moves more quickly, in opposition to the prediction arising from an association of causality with force.

Noisy Newton for Causal Inference

Michotte argued that if causal impressions were learned from the environment, these causal impressions should be related to the effect that one object has upon another. Later researchers have argued that the impression of causality comes from comparing the perceived motion to internal schemas of causal behavior, and these schemas could be the retrieved memories of past events (Dittrich & Lea, 1994; Rips, 2011; Weir, 1978; White & Milne, 1997). Unlike Michotte's connection of causal impression to the size of the effect on the impacted object, for causal schemas the impression of causality can be driven by the frequency of the event in the past.

The study of covariation and causality has provided separate insight into judgments of causality, with researchers investigating which probabilistic quantity best explains people's ratings of causality. Candidates have included the contingencies or the difference between the probability of an effect with and without a cause present (Jenkins & Ward, 1965; Rescorla, 1968), the causal power model (Cheng, 1997), and more recently a full Bayesian computation of the marginal probability of a causal relationship (Griffiths & Tenenbaum, 2009)

$$p(C|O) = \frac{p(O|C)p(C)}{p(O|C)p(C) + p(O|NC)p(NC)} \quad (6)$$

where O is the observed values, C is a causal relationship, and NC is a non-causal relationship.

The noisy Newton framework model follows the probabilistic approach of Griffiths and

Tenenbaum (2009), which as noted by Rips (2011) is a version of a schema model. It treats causal impression as inference about whether or not there is a causal relationship.

For causality judgments, the noisy Newton framework makes two kinds of predictions. The first is the qualitative prediction that impressions of causality based on collisions can be captured as probabilistic inference about whether an observation resulted from a collision or just the chance movement of two objects. The second is the quantitative predictions that the noisy Newton model makes about how different variables and different assumptions about the collision influence this inference. We tested these predictions in two experiments.

Physical Causal Impressions from Probabilistic Inference

To test the qualitative prediction in our first experiment, we investigated a simple collision between two objects with equal mass. In one condition, *real or random*, we instructed participants to make a decision in the way that the noisy Newton model decides: compare the hypothesis that the display was a real collision or came from an explicitly described non-collision alternative. In the second condition, *physical causality*, we instructed participants with more standard instructions: they were told to judge whether a display was causal or not, and were also given instructions that implied that they should treat the displays as real physical objects. The focus on physical causality, in a way that differed from instructions from previous causality experiments, was in order to focus on the mechanical types of causality that the noisy Newton model embodies. The qualitative prediction for this first experiment is that people make these sorts of judgments through a model comparison process and thus the results should be the same in the two instruction conditions.

Methods. Thirty participants from a university community were divided into groups of fifteen each and run in the two instruction conditions. Participants were seated so that their eyes were approximately 44 cm away from the display.

We indicated to all participants that they would be viewing a very simplified physical system by telling participants that the blocks were sliding along an invisible smooth surface and

that each of the blocks were made of the same material and had the same mass. Following these generic physical instructions, instructions specific to each condition were given to participants. In the *physical causality* condition, the additional instructions were that participants were to decide whether the gray block caused the white block to move or whether the white block moved by itself. After each movie, participants in this condition were asked, “Did it look like the white box moved because the gray box hit it? Was the white box’s movement produced by the gray box? Or did the white box take off on its own?”

In the *real or random* condition the additional instructions were,

Your task is to decide whether each movie came from a real collision of the blocks or a random combination of the variables. A real collision looks like the blocks actually collide. A random collision looks a little like a real collision, except that the velocities of the blocks, gap between the blocks, and the time delay before the second block starts moving are all selected randomly. Remember, both blocks always have the same mass.

Following these instructions, *real or random* participants were shown the boundaries of each of the variables. These instructions were meant to convince participants to use a uniform distribution over each of these variables as their random distribution. After each movie, participants responded by keypress as to whether the trial was a real collision or was random.

Three hundred trials were presented to each participant, with half drawn from a near-Newtonian collision distribution and half drawn from a distribution with heavier tails. The trials were drawn in this way to make the instructions seem plausible. The velocities ranged from 6 cm/s to 15 cm/s. The gap ranged from 0.1 mm to 4 mm, and the time delay ranged from 0 to 250 ms. All samples that fell outside the bounds of the variables were resampled. Each movie began with the fading in of a central white block and a gray block positioned 6.75 cm left of center.

Results and Discussion. The responses from the *real or random* condition and the *causality* condition were extremely similar, as shown in Figure 9. To quantify the results, we divided the gap, time delay and ratio of initial to final velocity into two bins each with approximately equal numbers of trials (using the values 0.04 cm, 0.065 s, and 1 as the respective

cutoffs for each variable). The bins were crossed between variables to yield eight cells and the percentage of trials judged to be a collision was computed for each cell for each participant. We investigated the results with a mixed effects ANOVA with condition as a between-subjects variable and the binarized velocity ratio, gap, and time delay as within-subject variables. There were main effects of gap, time delay, and velocity ratio, all $F(1, 28) > 30$, $p < 0.01$. There was a marginal interaction between velocity ratio and time delay, $F(1, 28) = 3.56$, $p = 0.07$, but no other interactions approached significance, $F(1, 28) < 1.4$, $p > 0.1$. The lack of effect of condition was suggestive and we further investigated the correspondence between the two conditions by computing the Pearson correlation between the eight cells for each condition averaged over participants. The correlation between the *real or random* condition and the *causality* condition was $r = 0.998$, showing that there was very good agreement between the conditions.

The similarity between conditions is consistent with participants' impression of causality arising from performing probabilistic inference to choose between the two categories. We cannot rule out the possibility that participants ignored the instructions in one of the conditions, but the similarity in judgments between our two conditions is striking as differences in instructions (using other instructions) have had large effects on responses in other causality judgment experiments (Schlottmann & Anderson, 1993; Schlottmann & Shanks, 1992), including in our second causality experiment.

Dependence of Physical Causal Impression on Velocities

The data from this experiment provide a test of the quantitative predictions of the noisy Newton model as well. We model the experiment as probabilistic inference about whether there was a causal relationship or if the display appears to be a result of random movement. As a simplification, we can assume that the probability of any stimulus arising from random movement is equal (though this depends on the scaling of the space), and therefore every term in Equation 6 is independent of the observed variables in the display except for $p(O|C)$. We can then compare

this quantity to an adjusted threshold T that accounts for the combined display-independent effects of $p(C)$, $p(O|NC)$, and $p(NC)$ to produce $p(C|O)$.

To maintain a clear link between the mass and causality judgment tasks, we reuse the same parameterized prior distributions from the mass judgment model and only introduce parameters for variables that were not relevant in the mass judgment task. The probability of the observed variables arising from a causal situation, $p(O|C)$, depends on the observed velocities, gap, and time delay. The prior probabilities of various gaps and time delays were irrelevant for the mass judgment task and need to be set here. If the bodies are perfectly rigid, then the smallest gap between them will be zero and neither will deform during the collision. Also, for rigid objects the time delay between the motor object stopping and the projectile object beginning to move approaches zero (Stronge, 2000). We will thus take a distribution with its entire mass on zero as our prior distribution for both time delay and gap, though we will discuss more realistic values later.

Despite having prior distributions of zero time delay and zero gap for causal interactions between variables, the noisy Newton model of causal judgments does not consider all non-zero gap and time-delay stimuli to be non-causal. The reason for this, as with the mass judgment model, is the noise in perception. A short observed time delay could be the result of noisy perception of a time delay of zero. Here we parameterize the noise for gap and time delay in the same way we did for velocities, as approximately following Weber's law, as is consistent with psychophysical studies (Andrews & Miller, 1977; Ekman, 1959). The noisy Newton model used the same velocity parameters as in the mass judgment task (σ^2 , w , and k_v) and three additional parameters (k_g , k_t , and p_r ; values given in Appendix A) were fit to produce the causality judgment task predictions. We show the graphical model for causal interactions in Figure 10, with the *real or random* and *physical causality* conditions combined together. One change we introduce is that the mass variables are tied together with a non-directional link – this is used to represent the assumption that the masses are equal, which participants in both conditions of this experiment were instructed to believe.

The proportion of collision judgments for different values of the experimentally manipulated variables are shown in Figure 11. This proportion decreases as the stimuli become less like physical collisions. In addition, we replicated the findings reported by Michotte (Experiments 39 and 40) and Natsoulas (1961): There is an asymmetry in the proportion of collision judgments such that when the final speed is less than the initial speed, participants believe the event is more like a collision. We evaluated the proportion of collision judgments for trials in which there was a large difference between the initial and final velocity, comparing judgments if the speed ratio was greater than 2/1 to judgments if the speed ratio was less than 1/2. If the initial velocity was twice as large as the final velocity or larger, the mean probability of collision judgments across participants was 0.14 higher compared to trials where the final velocity was twice as large as the initial velocity, a difference that was significant over participants, paired-sample $t(29) = 2.62, p < 0.05$.

With the same velocity parameters and prior as before, the noisy Newton model makes predictions that match the human data (Figure 11; NLL = 24,836 and AIC = 49,685). As a method for evaluating the performance of the model, we determined that the model predictions would match participant responses on 73% of trials. This was a strong result given the model produces noisy responses. As a comparison we predicted participant responses from participant responses: we divided the data into bins and predicted the causality responses in each bin from the mean number of causality responses in each bin (15 equally spaced bins per dimension as used in Figure 11, allowing a maximum total of 15^3 sets of data). This benchmark model matched participants responses on 73% of trials as well.

Despite assuming symmetric noise, the model produces asymmetries due to the assumption that collisions exhibit coefficients of restitution between zero and one. Using a value of $e = 1$ and with equal mass objects, the initial velocity of object a is equal to the final velocity of object b , but if $e < 1$, the initial velocity is greater than the final velocity. If an observer strictly believed that the masses were equal and $v_b = 0$, then this would imply that $e = 1$. As the model assumes a prior distribution over e that allows values less than 1 and noise in the velocity observations, it

predicts a positive skew in the choice function. The match between the model predictions and human data suggests that participants are sensitive to the range of the coefficient of restitution.

For causality judgments, the noisy Newton model predicted a 0.07 larger probability of choices in favor of collisions if the initial velocity was twice as large as the final velocity or larger compared to trials in which the final velocity was twice as large as the initial velocity². We performed a robustness analysis that included the same manipulations of the parameters as in the mass judgment task, but also manipulated the additional causality parameters, k_g , k_t , and p_r , by sequentially setting them at one-fourth, one-half, double, and quadruple their initial values. This resulted in twenty-eight additional sets of predictions. For 93% of the replications the asymmetry prediction was in the same direction as the human data.

Pure Causal Impressions from Probabilistic Inference

The first causality experiment gave evidence for both of our predictions: that qualitative prediction that impressions of physical causality can be captured as a probabilistic inference and the quantitative predictions that the noisy Newton model makes about how different variables. However, experiments on causal impression often attempt to avoid any implication that the display is of physical objects, instead asking participants to make causal judgments without assuming the objects are physical (e.g. Schlottmann & Anderson, 1993). To investigate whether the noisy Newton model can describe causality judgments with instructions of this type, a condition we term *pure causality*, we collected data from new participants under these instruction conditions. In addition, we replicated the *real or random* condition from the first causality experiment as a comparison. Because it was necessary to remove the physical constraint that the masses were equal from the *pure causality* condition, the noisy Newton model predicts a difference between conditions due to the physical constraints changing between the instruction conditions.

²A larger asymmetry in choice between can be produced by the noisy Newton model if the prior distribution over e is not uniform, but is given greater probability for smaller values of e .

Methods. Thirty participants from a university community were divided into groups of fifteen each and run in the two instruction conditions. The design was the same, except that the participants in the *pure causality* condition were only told to judge whether one square caused the other to move, that is was “a matter of subjective impression, there was no right or wrong answer, and no real collision was ever involved”.

Results and Discussion. The pattern of responses from the *real or random* condition and the *pure causality* condition were again similar, as shown in Figure 12 in the Data column. The Pearson correlation between conditions was $r = 0.98$. However, here the mixed effects ANOVA showed a main effect of condition, $F(1, 28) = 6.74$, $p < 0.05$ as well as main effects of gap, time delay, and velocity ratio, all $F(1, 28) > 19$, $p < 0.001$. Interactions between the variables were all insignificant or appeared to be the result of differences in random sampling of trials between the conditions³. As in the previous experiment, we tested of the asymmetry by comparing the judgments of causality for speed ratios that were greater than 2/1 to judgments if the speed ratio was less than 1/2. The difference was significant, paired-sample $t(29) = 3.74$, $p < 0.001$.

The results showed a similar pattern of dependence on the observable variables for both conditions, broadly consistent with the prediction that participants’ causal impressions are the result of probabilistic inference. However, in contrast to the first experiment, here there were real differences between the real or random and the pure causality condition. For these types of trials,

³The only significant interaction was between the instruction condition, velocity ratio, and time delay, $F(1, 28) = 10.1$, $p < 0.01$. We examined the three-way interaction by splitting the data into short and long time delays and running mixed effect ANOVAs on each of the restricted data sets. For the short time delay (less than 0.065s) there was no significant effect of condition or any significant interactions, all $F(1, 28) < 1.6$, $p > 0.1$. For the longer time delay, there was a significant main effects of condition, $F(1, 28) = 8.3$, $p < 0.01$, and a significant interaction between condition and velocity ratio, $F(1, 28) = 13.4$, $p < 0.01$. Surprisingly, the noisy Newton model produced this interaction between condition and velocity ratio exclusively for the short time delay using exactly the same model for both conditions (either masses assumed to be equal for both conditions or unequal for both conditions). As the noisy Newton model with the same assumptions in both conditions could not produce this interaction if the two conditions consisted of the same displays, it appears that this interaction is mainly an artifact of sampling the parameters of the displays.

the real or random instructions produced a stronger impression of causality than the pure causality instructions. A similar difference between conditions was found in Schlottmann and Anderson (1993) where the velocity ratio had a smaller impact on judgments of causality in the causality instruction condition, but there it was found that the effect of the gap was larger for the causality instruction condition which we did not replicate here⁴.

Noisy Newton Model Predictions. To address the differences between the instruction conditions in this experiment, we simulated the responses of the noisy Newton model using the same parameterized prior distributions as in the first experiment. The only change made to the models was a reflection of the physical constraints given in the instructions. Unlike in the first experiment, here the participants in the pure causality condition were not instructed that the masses were equal. To reflect this change, the masses were constrained to be equal for modeling the real or random condition, but were not constrained to be equal (using the same prior distributions as for the mass judgments) for the pure causality condition.

The predictions of the noisy Newton model are shown in the Model column of Figure 12. The predictions matched the results fairly closely (NLL = 25,097 and AIC = 50,205, equivalent to correctly predicting people's choices 72% of the time compared to 73% for the benchmark model as described in the first causality experiment). First the two instruction conditions showed a broadly similar dependence of the impression of causality on the variables of the display, with a predicted asymmetry of 0.04. More interestingly, the noisy Newton model produced differences between conditions quite similar to the differences found in the participants' responses, with the real or random condition showing stronger impressions of causality than the pure causality condition. The predictions were robust to different parameters, with a velocity asymmetry for 55% of the parameter sets and with the real or random condition producing larger causality judgments for 93% of the parameter sets tested.

The main effect of instruction condition is both mechanically and probabilistically driven.

⁴A possible explanation for this advanced by a reviewer was that the larger velocities used in the Schlottmann and Anderson (1993) study produced larger causal impressions that made smaller effects more apparent.

If the masses are not constrained to be equal, then more combinations of initial and final velocity are possible. Given that the motor object stops as a result of impacting the projectile object, we can manipulate the equations of Newtonian mechanics to show the final velocity of the projectile object is equal to the coefficient of restitution multiplied by the velocity of the motor object: $v_b = eu_a$. As the coefficient of restitution is constrained to lie between zero and one, any final velocity less than the initial velocity is physically possible if the coefficient of restitution is believable. The reduction in causal impression comes from treating the belief in coefficient of restitution as a probability distribution: for a sharp prior when the masses are constrained to be equal, the probability for a collision with equal initial and final velocities will be higher than for a broad prior for which this combination is still possible.

Michotte (1963) took the result that the causal impression was stronger when the projectile object moved less as an indication that causal perception and mechanical causality diverged, but our analysis shows how it can result from inference using noisy perception and Newtonian mechanics. Overall, the results shown here are evidence that the noisy Newton model can produce results that have been taken to dissociate causality judgments from experience in a world well described by Newtonian mechanics. However, we cannot conclude from only these results that causal impression is dependent only on physical hypotheses as there remain other effects that need to be explained, such as higher velocities leading to stronger impressions of causality (Michotte, 1963), and aspects of the dependence of the impression of causality on time (as discussed below in the Future Directions section).

Comparison to Other Models of Causality

There are of course other models of causal impression. Here we compare the noisy Newton model to two other models of causality judgments: the model of L. B. Cohen, Chaput, and Cashon (2002, CC&C) and the model of Schlottmann and Anderson (1993, S&A). We do so qualitatively because while both these models could be adapted or fit to likely provide good quantitative descriptions of the data, the steps necessary to provide a good quantitative fit

highlight the similarities and differences with the noisy Newton model.

The CC&C model is an artificial neural network that uses the ideas of Hebb (1949) and Kohonen (1997) to learn causal impression from experience with collisions. The model, constructed to produce similarity relationships found in children (L. B. Cohen & Amsel, 1998; Leslie, 1984), consists of hierarchically organized layers: the two bottom layers receive input exclusively from either movement or position information respectively, while the top layer receives input from both of the bottom layers. The activation of the top layer of the CC&C model depends on how well the current stimulus matches the training stimuli, because of this it would make predictions similar to that of the noisy Newton model if the training stimuli follow Newtonian mechanics. A difference between the two is that the training trials the CC&C model is exposed to are combined into the weights rather than having an explicitly represented prior as the noisy Newton model does. The explicit representation of the noisy Newton model makes it possible to isolate the regions of the prior distribution consistent with instructions and thus predict a difference between equal mass and unequal mass instructions. The CC&C does not have a mechanism that would allow it to predict this difference.

In contrast to the CC&C and the noisy Newton model, the S&A model does not specify how causal impression depends on the physical variables of the stimuli. This relationship is instead inferred through fitting the model to the responses made by participants. The S&A model is able to make this inference because it is an account of how causal impression is combined across physical dimensions: Given the contribution to causal impression for each setting of each of the physical variables and the weights associated with each setting of each physical variable, a prediction can be made for the overall causal impression. This model can be fit and evaluated on factorial designs that provide many combinations of levels, but the data from our two causality experiments are especially inhospitable for this model as the physical parameters for each trial were chosen from continuous distributions, leaving little overlap in parameters between trials. As a result of this design, the straightforward application of the S&A model results in a perfect fit but a tremendous penalty for the number of parameters ($NLL = 0$ and $AIC = 105,600$ for the first

experiment and $NLL = 0$ and $AIC = 108,000$ for the second experiment). In order for the model to be uniquely specified for these experiments, the data need to be binned in some fashion and here the S&A model's match to the human data are bounded by the benchmark model, as it is essentially a more constrained version of this benchmark. The strength of the S&A model is that it can explain interactions between settings of physical variables, but effects such as the asymmetry in response for high and low speed ratios and the influence of instructions cannot be predicted a priori by this model.

Learning Newtonian Mechanics

For both mass judgments and causality judgments, the noisy Newton framework depends on an accurate conception of Newtonian mechanics. If this assumption is accepted, it naturally leads to the question of whether such an accurate conception could be learned, or whether it must be innate. This is a particularly relevant question for the study of the impression of causality, in which one of the biggest controversies is whether this impression is innate or whether it could be learned through experience. Michotte (1963) made the strong claim that this impression was innate and that observers were nearly unanimous in their judgments. Unanimity of response is an argument for innateness, but later researchers found many individual differences in whether a display delivered a causal impression and in the strength of the causal impression (Beasley, 1968; Boyle, 1960; Gemelli & Cappellini, 1958; Schlottmann & Anderson, 1993; Straube & Chatterjee, 2010). Even now, the question of innateness still has not been settled (Rips, 2011; Saxe & Carey, 2006; Scholl & Tremoulet, 2000), because even developmental evidence that infants are sensitive to causality from as young as six months of age (Leslie, 1984; Leslie & Keeble, 1987) could be due to innate mechanisms or due to learning.

Here we show that it is possible to learn to make judgments based on probabilistic inferences that approximate Bayesian inference in our model via a simple algorithmic process that does not require direct evaluation of the equations of Newtonian mechanics. The idea is to use the memories of people's experience with colliding objects in the world — which are

guaranteed to be consistent with Newtonian mechanics — to stand in place of the noisy Newton model. We make two assumptions. First, we assume that people can store a small number of past experiences with collisions. As Newtonian mechanics is a good description of each of these collisions, by assuming the experiences chosen are drawn randomly they will be representative of the environmental prior. Some of these memories would need to include mass information as well as velocities, but experience could give access to this information as well. Second, we assume that people can assess the perceptual similarity of a given collision with each of these stored memories, using a similarity function derived from perceptual noise.

Given these capabilities, we can approximate Bayesian inference in the noisy Newton model by comparing the summed similarity of a new event to memories of previous events. For mass judgments, the summed similarity of the new event is compared to memories in which $m_a > m_b$ and to memories in which $m_b > m_a$, and the response with the greater summed similarity is more often chosen (as in A. Cohen, 2006; Lamberts, 2004; Nosofsky, 1986). For causality judgments, the summed similarity of the new event is compared to a threshold. Both of these processes are very much the idea of a schema in the sense of Rips (2011), and with this approximation people would be merely storing the results of the environmental computation and using knowledge of their perceptual noise to compute similarity – no explicit model is inferred.

In fact, this strategy has been shown to provide a general method for approximating simple Bayesian inferences (Shi, Griffiths, Feldman, & Sanborn, 2010). Instead of directly implementing Bayes' rule applied to a precise mathematical description of the Newtonian model with precise priors, memories (by definition, drawn from and consistent with the environmental prior) are compared to the current event, according to the likelihood model induced by perceptual noise. More formally, we can re-express the posterior probability of a hidden variable setting S given observed variables O as:

$$P(S|O) = \frac{P(S)P(O|S)}{P(O)} \approx \frac{P(O|S)}{\sum_{S'} P(O|S')} \quad (7)$$

where the sum in the denominator on the right hand side ranges over all instances, S' , of the

hidden variable stored in memory. Here $P(O|S)$ is zero if S is not stored in memory.

To see if the learning process is plausible, we need to show that behavior that matches the full prior can be generated from experience with a small number of events that were not perfectly perceived. To avoid accuracy issues due to stored experience in the regime where the model's judgments are biased, we only use stored samples where one mass is much larger than the other, collision events for which human judgment is near perfect (Runeson et al., 2000).

We modified the prior distribution over the masses so that there is an equal probability that $m_a \sim \text{Exponential}(1) + 1000$ and $m_b \sim \text{Exponential}(1)$ or $m_a \sim \text{Exponential}(1)$ and $m_b \sim \text{Exponential}(1) + 1000$. In these cases the ratio m_a/m_b is either very large or very small and it would be obvious which object is heavier. The modified mass distributions and the prior distributions used before for the other variables were taken as the prior distribution of the events in the external environment. Fifty samples from the prior were drawn for each simulated participant and each sample was corrupted by perceptual noise. The noisy samples from the prior distribution were used to make predictions for that simulated participant, and many simulated participants were averaged together to produce the mean prediction. The predictions of the physical model with these priors are shown in Figure 13, using the same parameters as before for the mass judgments and only modifying the threshold for the causality judgments⁵. The predictions for both mass judgments and causality judgments are very similar to the predictions of the original model (see Figures 5, 7, and 11 for comparison).

Special care needs to be taken to explain how this process could be implemented by infants. Our own implementation of causal judgments in the noisy Newton framework compared the probability that an observed event was a collision to a threshold. Applying this implementation to the dishabituation experimental designs usually used in studies of causality with infants would mean that infants can judge whether displays are above or below threshold and are surprised if a series of below threshold events is followed by one above threshold. More generally, non-causal hypotheses could take forms other than randomness, such as the hypothesis of tunneling that we

⁵The threshold was weakened from 10^{-10} to 10^{-15} .

discuss in the General Discussion. Learning that events in the world should be described as different hypotheses does not necessarily require innate knowledge, instead infants could form latent clusters of different types of objects, such as those that move mechanistically and those that move under their own power. Assuming these events do cluster together, latent clustering of this sort of statistical structure is a route toward predicting the infant dishabituation results.

This mechanism for learning Newtonian mechanics is similar to other previous proposals for learning intuitive physics. A. Cohen (2006) proposed an exemplar mechanism for learning, which was able to successfully predict the effect of training particular collision events on the mass judgments of similar collision events. These training trials influenced participants' later judgments for stimuli with similar physical parameters, but not for stimuli with very different physical parameters. Our framework differs from that of A. Cohen (2006) in that the knowledge that participants bring into the experiment matches Newtonian mechanics, rather than using a set of heuristics as the prior knowledge. Relatedly, Friedman and Forbus (2009) (see also Friedman, Taylor, & Forbus, 2009) proposed a sophisticated approach for learning from exemplars within the qualitative physics framework. The noisy Newton model differs from this approach because it represents variables quantitatively.

General Discussion

We have developed a framework for explaining how we infer different aspects of the physical world from noisy sensory information, how these inferences are underpinned by an accurate knowledge of Newtonian mechanics, and suggested that the knowledge of Newtonian mechanics could be learned by simply remembering previous experiences with colliding objects. This framework can be easily extended to a variety of judgments using different types of stimuli – all that is required is that we characterize the sensory noise as well as how the stimuli should behave using Newtonian mechanics.

The components of the noisy Newton framework combined to predict a wide variety of empirical effects of mass and causality judgments. Some predictions were straightforward:

adding sensory noise to Newtonian mechanics predicted the effect of the coefficient of restitution on accuracy of mass judgments and the asymmetric effects of velocity ratio in causal impression. Other predictions were more surprising. The motor object bias was predicted by an interaction of Newtonian mechanics with a reasonable prior belief about velocities, and changes to the prior and changes to the likelihood predicted the effects of training and occlusion on the motor object bias respectively. In addition, the differences between the instruction conditions in the causality experiment were predicted by a loosening of the constraints on the physical situation – this spread the prior distribution out, reducing the predicted causal impression of the noisy Newton model.

In this discussion, we review results using other hypotheses, other cues, and other types of judgment that show a general qualitative agreement with our approach, examine what kind of intuitive knowledge do we possess, and look at future directions for the noisy Newton framework.

Other Effects Using Colliding Objects

We have shown quantitatively that our model agrees with a series of results that have been taken as evidence of dissociations between intuitive and Newtonian mechanics. Here we look at wider applications of our method and how it qualitatively predicts results from a broad range of studies using Michotte-like stimuli, examining how we could use additional hypotheses, model additional observable variables, and explain additional types of judgment.

Additional Hypotheses. In our application of the noisy Newton framework to causality above, we considered only the possibility of a causal relationship that followed Newtonian mechanics and a simplified non-causal relationship based on the random movement of the stimuli. Michotte (1963) allowed a much broader range of responses, including such causal impressions as one object carrying off the other (entraining) or one object setting off the other object (triggering).

The causality framework we used is extensible to these other impressions, though they do not all correspond to simple Newtonian mechanics. Causal impressions of triggering depend on an energy source within the stationary object, not just momentum transferred from the motor object. Heider and Simmel (1944) demonstrated many other examples of simple objects that gave

people the impression of intentions, moods, and even personality traits. Building an appropriate prior distribution over these movements could allow our model to infer what generated the behavior, bringing the model more in line with causal schema approaches such as that of Weir (1978). This sort of model would need to marry a wider conception of physics with descriptions of behavior.

One impression that does seem to have a mechanical basis is that of tunneling, which is the impression that the motor object passes by the projectile object and the projectile object remains stationary. In Michotte's Experiments 7-10, it was shown that small manipulations that made the display harder to see changed the basic launching impression into one of tunneling. If participants fixated at a point below the point of impact (Experiment 7), or if a semi-transparent sheet of paper made the objects less distinct to participants (Experiment 8), or if participants were moved from 1.5 to six or seven meters away from the apparatus (Experiment 9), or even if very tiny stimuli were used (Experiment 10), the causal impression changed from launching to tunneling.

These results could be explained in the noisy Newton framework as a choice between mechanical causal hypotheses. The stimuli Michotte used in these experiments tend to produce the causal impression of launching, but if additional noise is introduced in perception, the causal impression becomes one of tunneling. A higher prior probability of tunneling rather than launching could produce this effect in the noisy Newton framework, so that when the evidence is weak the prior bias toward objects missing each other would produce the causal impression of tunneling. This prior bias is environmentally plausible because if there is reasonable uncertainty about the depth of the objects in the display and the objects are not too large, it is more likely that the objects miss each other rather than collide. This hypothesis is bolstered by Michotte's Experiment 37, in which the launching impression was destroyed when the two objects were projected onto screens at obviously different depths.

Additional Cues. The noisy Newton framework can also be extended to accommodate extra sources of information. Because we assume our observed variables are noisy, there is always room for additional variables to provide more information about what state generated the

observed variables. For mass judgments, one obvious piece of information is the difference in size of objects. The influence of size on relative mass judgments was studied by Natsoulas (1960) who found that the larger object of a pair was judged to be the heavier object more often, as would be expected from our experience with the environment.

Like mass judgments, the impression of causality also can be influenced by additional cues, such as a sound that occurs at the time of the collision. Sekuler, Sekuler, and Lau (1997) investigated a display in which two identical objects moved toward each other, overlapped, and then moved away from each other. The display was ambiguous as to whether the objects had simply moved past one another or had collided and reversed direction. Consistent with the prior belief posited above that two objects tend to be at different depths and would miss each other, the perception of the objects moving past one another dominated a purely visual display. However, it was shown that a sound played at or near the time of collision increased the causal impression, while a constant sound that disappeared only at the time of collision did not have an effect.

Scholl and Nakayama (2002) showed how two spatially and temporally neighboring events could influence one another. A display may appear to be a tunneling event alone, but when paired with a nearby launching event, the perception of tunneling changes to one of launching. This effect was termed *causal capture*, and can even be used to change the perception of what happens behind an occluder (Bae & Flombaum, 2011). We can hazard an explanation for this result from the noisy Newton framework as well, by assuming that participants use motion cues to group the display of four objects into two pairs of two objects (e.g., Hofsten & Spelke, 1985). Then the clear launching event for one pair will influence the overall impression of what is happening between the larger groups. Consistent with this, Choi and Scholl (2004) demonstrated how perceptual grouping was important for either producing causal capture or not, as a nearby launching event moving in same direction as the target objects producing a larger causal capture effect compared to the nearby launching event moving in the opposite direction as the target event.

Additional Types of Judgment. We have presented evidence that the noisy Newton framework is not limited to one type of judgment as the heuristic model is, but can predict both

people's judgments of relative mass and impression of causality. Here we review research on people's perceptions of velocity, time delay, gap, the coefficient of restitution, and even force with colliding objects to show that they are often qualitatively consistent with the noisy Newton framework.

Parovel and Casco (2006) carefully investigated the relationship between the causal impressions of launching and triggering and the perceived velocity of the projectile object, using a task in which participants compared the projectile object's velocity to a standard moving object. When the velocity of the motor object was greater than the velocity of the projectile object, participants were biased to believe that the velocity of the projectile object was greater than an equally speedy projectile object in a non-causal display. Speedier motor objects produced larger overestimates of the speed of the projectile object, indicating that participants had integrated the two velocities when participants reported a causal impression of launching. Crucially, the overestimate of the projectile object's velocity did not depend on the velocity of the motor object when the projectile object moved faster than the motor object. This showed that the integration of the velocities was not purely a low-level averaging. Instead, when participants reported a causal impression of triggering the estimate depended only on the velocity of the projectile object. If estimating the projectile object velocity using the noisy Newton framework, we would expect the motor object velocity to influence our judgment: The motor object velocity is informative about the expected projectile object velocity. However, for displays which gave the causal impression of triggering, if an object activates its own motion in response to the presence of another object, its velocity should not depend on the speed of the first object.

Time delay perception has been investigated as well, using the presence or absence of a sound (Guski & Troje, 2003). As expected from adding another source of information that indicates a collision, the presence of a *clack* sound decreased participants' perceived duration of the time delay.

Judgments of the gap or the overlap between objects during collision can be influenced by the causal capture effect (Scholl & Nakayama, 2004). Displays were presented in which one set

of objects partially or completely overlapped and these displays alone gave the causal impression of tunneling. However, if an unambiguous launching event was also presented nearby, participants underestimated the amount that the objects overlapped. This is consistent with participants having a prior belief that colliding objects do not overlap too much, and this prior belief resulting in a reduction of the perceived overlap when there is contextual evidence that a launching event occurred.

Like the judgments of mass which were about unobserved variables, people can also make judgments about the coefficient of restitution of objects. Warren, Kim, and Husney (1987) investigated participants' perception of the "bounciness" of balls and attempted to relate the ratings made to the coefficient of restitution of the objects. There was a high correlation between bounciness and the coefficient of restitution, even if some of the information used in the judgments was occluded. Given this correlation, we could model bounciness judgments within the noisy Newton framework as estimates of the coefficient of restitution, but would have to account for how some physical variables appear more correlated with bounciness judgments and motor control than others (Siegler, Bardy, & Warren, 2010; Warren et al., 1987).

Not all judgments seem in accord with the noisy Newton framework; judgments of force appear to be the exception. Newton's third law clearly describes how force applies to each object in a collision: Both objects exert the same amount of force on each other. A dissociation between force and causality is that participants do not report that the projectile object stopped the motor object, only that the motor object set the projectile object into motion (Michotte, 1963; White, 2006). Investigating people's judgment of force alone, White (2009) described how the judgments people make about forces are divided into force and resistance, depending on which object is perceived to act upon the other. The split of force into two roles could be due to a dissociation of the common meaning of force from the meaning used by physicists (Talmy, 1988), but the difference in strengths is a more fundamental deviation of force judgments from physical force (White, 2006). It remains to be seen whether this is a dissociation from Newtonian mechanics or whether a remapping of variables of the Newtonian framework could correspond to people's

intuitive ideas about force.

What Sort of Knowledge of Mechanics Do We Have?

The power of physical theories, be they Aristotelian, medieval impetus, Newtonian, or modern, is that they attempt to provide a consistent explanation for all physical phenomena. People, however, do not seem to display such consistency. Researchers have proposed that our understanding of mechanics reflects formal pre-Newtonian systems of ideas, such as Aristotelian (diSessa, 1982; Shanon, 1976) or medieval impetus theories (Hubbard & Ruppel, 2002; Kozhevnikov & Hegarty, 2001; McCloskey, 1983) of mechanics. However, these proposals have had difficulty explaining how people's responses appear to reflect different systems of ideas depending on the particular problem that they have been given (Ranney & Thagard, 1988).

The puzzles that we referenced at the beginning of the paper demonstrate this inconsistency as well. While people are poor at reasoning about a ball that has followed a curved path (McCloskey et al., 1980), they are much more accurate at reasoning about water that has passed through a curved hose (Kaiser, Jonides, & Alexander, 1986) or about videos of balls traveling through tubes (Kaiser, Proffitt, & Anderson, 1985). Likewise, while participants predict that a ball dropped from a moving object will fall straight downwards (McCloskey et al., 1983), when shown animated balls dropping from moving objects, they tend to pick the object moving forward as the correct option (Kaiser et al., 1992).

A proposed explanation for our biases and how they change with observing animations is that people can only reason about one-dimensional quantities (Proffitt & Gilden, 1989). We can extract certain pieces of information from displays, but when the decision requires us to use multiple sources of information, we are unable to combine them effectively. Animation was described as allowing participants to temporally segregate useful one-dimensional information, such as the path of the ball when it leaves the mouth of the C-shaped tube or the location of the dropped ball relative to the moving object that dropped it (Kaiser et al., 1992). Mass judgments have also been used to justify the one-dimensional approach, because of the bias found in

people's judgments. We have shown here that the mass judgment bias can be better explained by correct application of Newtonian mechanics in the presence of sensory noise. In addition, Hamrick et al. (2011) have used a noisy Newton approach to explain assessments of the stability of 3D objects, a problem in which the objects cannot be treated as one-dimensional point masses.

To explain these dissociations in reasoning in the noisy Newton framework may require us to place limits on the generality of the prior information. A strong version of the noisy Newton framework would hold that participants have perfect prior knowledge in all situations, and this prior knowledge is projected onto the variables of interest. However, we could also have local knowledge of Newtonian mechanics that corresponds to our perceptual variables, but not be able to project it onto other representations of the problem. This knowledge could be innate or learned, but would predict that our accuracy in these tasks would diminish as we distance ourselves from the variables with which we are familiar. This echoes the weak KSD approach, in which people only have an accurate conception of Newtonian mechanics within situations that match our terrestrial experiences (Gilden, 1991; Warren et al., 1987).

Future Directions

There are several interesting directions in which to expand the noisy Newton framework. Staying within the domain of judgments from simple collisions, an obvious direction is to test what the model predicts for either training experiments or experiments in which information is occluded. For training experiments, predictions can be made about the time course and final result of training. For occlusion experiments, removing information could allow for predictions from the noisy Newton model to be compared with predictions from heuristic models when the specific information required for a heuristic is removed (e.g., Gilden & Proffitt, 1989; Runeson & Vedeler, 1993).

A second direction is to investigate how well more realistic models of mechanics correspond with human judgments. Both Schlottmann and Anderson (1993) and Michotte (1963) found that collisions with delays of tens of milliseconds produced the best causal impressions.

Our data appear to show the same pattern, as the percentage of causal impression for the 5 ms time delay was significantly less than for the 15 ms time delay, $t(1846) = 2.70$, $p < 0.01$. The classical approximation that we used in our implementation for causality impressions assumed that the duration of contact between objects was zero, while the forces that act between objects are not instantaneous. One explanation is that people derive causal impressions when events occur together in iconic memory (White, 1988). An explanation using the noisy Newton framework would need to tie contact durations of tens of milliseconds to real-world stimuli that remain in contact that long. Collisions between metal balls or between a metal golf club and golf ball tend to be less than a millisecond (Goldsmith, 2001; Roberts, Jones, & Rothberg, 2001), but other collisions do last much longer. The duration of contact between the human head and a soccer ball is around 10 milliseconds (Rezaei, Verhelst, Van Paepegem, & Degrieck, 2011). Contact durations that very commonly experienced by participants are on the order of tens of milliseconds: around 30 ms between a finger and a light switch (Maij, Grave, Brenner, & Smeets, 2011) and around 100-150 ms for a well-practiced keypress (Kello, Beltz, Holden, & Van Orden, 2007). It is worth noting that detailed dynamical simulations — such as the rigid body dynamics simulators used in modern, physically realistic video games — maintain internal quantities that treat the process by which a collision unfolds in detail. The noisy Newton model can be extended to more realistic collisions by using these simulators to generate the prior distributions, rather than only using the closed-form solution we have studied here.

Going beyond simple collisions, we anticipate that the combination of Newtonian physics and Bayesian inference that underlies our framework will be useful in accounting for people's inferences in a wide range of settings. People have to make inferences about the properties of objects based on the observed behavior of those objects as a basic component of planning and executing motor actions, suggesting that the scope of noisy Newtonian models must be wider than just the simple collisions that we have focused on in this paper. Others have recently presented results suggesting that people may be able to stochastically simulate physical events for more complex sets of objects, providing a forward model of physical causality that could support

inferences about unobserved variables in a way that is very similar to our noisy Newton framework (Gerstenberg et al., 2012; Hamrick et al., 2011; Smith & Vul, 2012). Determining the assumptions that underlie people's expectations about physical events is also a key step towards understanding the circumstances in which people might postulate the presence of particular forces, whether they reflect hidden causes (Griffiths & Tenenbaum, 2009) or provide evidence for animacy (Tremoulet & Feldman, 2000).

Conclusions

Our results show that it is possible to explain people's inferences about collisions between objects in a way that is consistent with Newtonian physics, provided we use Bayesian inference to take into account the possibility of noise in people's observations. The resulting rational model of intuitive dynamics explains previous results that have been taken as evidence that people reason about collisions by applying simple heuristics. It also goes beyond these heuristic accounts in providing a way to understand how people might derive such a model from experience, and hence be helped by training, make decisions if information is missing, and by providing a single framework in which analyses of relative mass judgments can be unified with experiments on perceptual causality and other judgments as well.

The discovery of Newtonian physics was a major intellectual achievement, and its principles remain difficult for people to learn explicitly; in this sense, Newtonian physics and intuitive physics might seem far apart. On the other hand, over the course of cognitive development, people learn how to deftly interact with and reason about a physical world that is often well described in Newtonian terms, so our unconscious physical intuitions must be in some harmony with physical law. Combining Newtonian physics with Bayesian inference, explaining apparent deviations from precise physical law by the uncertainty in inherently ambiguous sensory data, thus seems a particularly apt way to explore the foundations of people's physical intuitions.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & F. Csaki (Eds.), *Proceedings of the second international symposium on information theory*. (p. 267-281). Budapest: Akademiai Kiado.
- Akaike, H. (1978). On the likelihood of a time series model. *The Statistician*, 27, 217-235.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3), 409-429.
- Andersson, I. E. K., & Runeson, S. (2008). Realism of confidence, modes of apprehension, and variable-use in visual discrimination of relative mass. *Ecological Psychology*, 20(1), 37-41.
- Andrews, D. P., & Miller, D. T. (1977). Acuity for spatial separation as a function of stimulus size. *Vision Research*, 18, 615-619.
- Bae, G. Y., & Flombaum, J. I. (2011). Amodal causal capture in the tunnel effect. *Perception*, 40(1), 74-90.
- Beasley, N. A. (1968). The extent of individual differences in the perception of causality. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 22(5), 399-407.
- Boyle, D. G. (1960). A contribution to the study of phenomenal causation. *Quarterly Journal of Experimental Psychology*, 12(3), 171-179.
- Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference: a practical information-theoretic approach*. New York: Springer-Verlag.
- Cheng, P. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367-405.
- Choi, H., & Scholl, B. J. (2004). Effects of grouping and attention on the perception of causality. *Perception and Psychophysics*, 66(6), 926-942.
- Cohen, A. (2006). Contributions of invariants, heuristics, and exemplars to the visual perception of relative mass. *Journal of Experimental Psychology: Human Perception and*

Performance, 32, 574-598.

Cohen, A., & Ross, M. (2009). Exploring mass perception with Markov chain Monte Carlo.

Journal of Experimental Psychology: Human Perception and Performance, 35, 1833-1844.

Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of the causality of a simple event. *Infant Behavior and Development*, 21, 713-731.

Cohen, L. B., Chaput, H. H., & Cashon, C. H. (2002). A constructivist model of infant cognition. *Cognitive Development*, 17, 1323-1343.

Cohen, L. B., & Oakes, L. M. (1993). How infants perceive a simple causal event. *Developmental Psychology*, 29, 421-433.

de Vries, H. L. (1943). The quantum character of light and its bearing upon threshold of vision, the differential sensitivity and visual acuity of the eye. *Physica*, 10, 553-564.

diSessa, A. A. (1982). Unlearning Aristotelian physics: A study of knowledge-based learning. *Cognitive Science*, 6, 37-75.

Dittrich, W. H., & Lea, S. E. (1994). Visual perception of intentional motion. *Perception*, 23(3), 253-268.

Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Boca Raton, Florida: Chapman & Hall / CRC.

Ekman, G. (1959). Weber's law and related functions. *The Journal of Psychology: Interdisciplinary and Applied*, 47, 343-352.

Flynn, S. B. (1994). The perception of relative mass in physical collisions. *Ecological Psychology*, 6(3), 185-204.

Friedman, S., & Forbus, K. D. (2009). Learning naïve physics models and misconceptions. In *Proceedings of the 31st annual conference of the cognitive science society*. (p. 2505-2510).

Friedman, S., Taylor, J., & Forbus, K. D. (2009). Learning naïve physics models by analogical generalization. In *Proceedings of the 2nd international analogy conference*.

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, 96, 267-314.

- Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, *41*, 711-724.
- Gemelli, A., & Cappellini, A. (1958). The influence of the subject's attitude in perception. *Acta Psychologica*, *14*, 12-23.
- Gerstenberg, T., Goodman, N., Lagnado, D. A., & Tenenbaum, J. B. (2012). Noisy Newtons: Unifying process and dependency accounts of causal attribution. In *Proceedings of the 34th annual conference of the cognitive science society* (p. 378-383).
- Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston, MA: Houghton Mifflin.
- Gilden, D. L. (1991). On the origins of dynamical awareness. *Psychological Review*, *98*, 554-568.
- Gilden, D. L., & Proffitt, D. R. (1989). Understanding collision dynamics. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 372-383.
- Gilden, D. L., & Proffitt, D. R. (1994). Heuristic judgment of mass ratio in two-body collisions. *Perception & Psychophysics*, *56*, 708-720.
- Goldsmith, W. (2001). *Impact: The theory and physical behaviour of colliding solids*. Mineola, New York: Dover.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*(4), 661-716.
- Guski, R., & Troje, N. F. (2003). Audiovisual phenomenal causality. *Perception & Psychophysics*, *65*(5), 789-800.
- Hamrick, J., Battaglia, P., & Tenenbaum, J. B. (2011). Internal physics models guide probabilistic judgments about object dynamics. In *Proceedings of the 33rd annual conference of the cognitive science society* (p. 1545-1550).
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York:

Wiley.

Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, *57*(2), 243–259.

Hick, W. E. (1950). The threshold for sudden changes in the velocity of a seen object. *The Quarterly Journal of Experimental Psychology*, *2*, 33-41.

Hofsten, C. von, & Spelke, E. S. (1985). Object perception and object-directed reaching in infancy. *Journal of Experimental Psychology: General*, *114*, 198-212.

Hubbard, T. L., & Ruppel, S. E. (2002). A possible role of naïve impetus in Michotte's "launching effect": Evidence from representational momentum. *Visual Cognition*, *9*(1-2), 153-176.

Hume, D. (1748). *An enquiry concerning human understanding*. Indianapolis, IN: Hackett.

Jacobs, D. M., Michaels, C. F., & Runeson, S. (2000). Learning to perceive the relative mass of colliding balls: the effects of ratio scaling and feedback. *Perception & Psychophysics*, *62*, 1332-1340.

Jacobs, D. M., Runeson, S., & Michaels, C. F. (2001). Learning to visually perceive the relative mass of colliding balls in globally and locally constrained task ecologies. *Journal of Experimental Psychology: Human Perception and Performance*, *27*, 1019-1038.

Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs*, *79*.

Kaiser, M. K., Jonides, J., & Alexander, J. (1986). Intuitive reasoning about abstract and familiar physics problems. *Memory & Cognition*, *14*, 308-312.

Kaiser, M. K., & Proffitt, D. (1984). The development of sensitivity to dynamically relevant causal information. *Child Development*, *55*, 1614-1624.

Kaiser, M. K., Proffitt, D. R., & Anderson, K. (1985). Judgments of natural and anomalous trajectories in the presence and absence of motion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*(4), 795-803.

Kaiser, M. K., Proffitt, D. R., Whelan, S. M., & Hecht, H. (1992). Influence of animation on

- dynamical judgments. *Journal of Experimental Psychology: Human Perception and Performance*, 18(3), 669-689.
- Kello, C. T., Beltz, B. C., Holden, J. G., & Van Orden, G. C. (2007). The emergent coordination of cognitive function. *Journal of Experimental Psychology: General*, 136, 551-568.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271-304.
- Kohonen, T. (1997). *Self-organizing maps*. Berlin: Springer.
- Kozhevnikov, M., & Hegarty, M. (2001). Impetus beliefs as default heuristics: dissociation between explicit and implicit knowledge about motion. *Psychonomic Bulletin & Review*, 8(3), 439-453.
- Lamberts, K. (2004). An exemplar model for perceptual categorization of events. *The Psychology of Learning and Motivation*, 44, 227-260.
- Leslie, A. M. (1984). Spatiotemporal continuity and the perception of causality in infants. *Perception*, 13, 287-305.
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25, 265-288.
- Maij, F., Grave, D. D. J. de, Brenner, E., & Smeets, J. B. J. (2011). Misjudging where you felt a light switch in a dark room. *Experimental Brain Research*, 213, 223-227.
- Marr, D. (1982). *Vision*. San Francisco, CA: W. H. Freeman.
- McCloskey, M. (1983). Intuitive physics. *Scientific American*, 284, 114-123.
- McCloskey, M., Caramazza, A., & Green, B. (1980). Curvilinear motion in the absence of external forces: Naive beliefs about the motion of objects. *Science*, 210, 1139-1141.
- McCloskey, M., Washburn, A., & Felch, L. (1983). Intuitive physics: the straight-down belief and its origin. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(4), 636-649.
- McIntyre, J., Zago, M., Berthoz, A., & Lacquaniti, F. (2001). Does the brain model Newton's laws? *Nature Neuroscience*, 4, 693-694.

- Michotte, A. (1963). *The perception of causality*. New York: Basic Books.
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, *434*, 387-391.
- Natsoulas, T. (1960). Judgments of velocity and weight in a causal situation. *The American Journal of Psychology*, *73*(3), 404–410.
- Natsoulas, T. (1961). Principles of momentum and kinetic energy in the perception of causality. *The American Journal of Psychology*, *74*(3), 394-402.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39-57.
- Notterman, J. M., & Page, D. E. (1957). Weber's law and the difference threshold for the velocity of a seen object. , *126*, 652.
- Parovel, G., & Casco, C. (2006). The psychophysical law of speed estimation in Michotte's causal events. *Vision Research*, *46*(24), 4134-4142.
- Peterson, W. W., Birdsall, T. G., & Fox, W. C. (1954). The theory of signal detectability. *Transactions of the Institute of Radio Engineers, Professional Group on Information Theory*, *4*, 171-212.
- Proffitt, D. R., & Gilden, D. (1989). Understanding natural dynamics. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(2), 384-393.
- Ranney, M., & Thagard, P. (1988). Explanatory coherence and belief revision in naive physics. In *Proceedings of the tenth annual conference of the cognitive science society* (p. 426-432). Hillsdale, NJ: Erlbaum.
- Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*, 1-5.
- Rezaei, A., Verhelst, R., Van Paepegem, W., & Degrieck, J. (2011). Finite element modelling and experimental study of oblique soccer ball bounce. *Journal of Sports Sciences*, *29*(11), 1201-1213.
- Rips, L. J. (2011). Causation from perception. *Perspectives on Psychological Science*, *6*(1),

77-97.

- Roberts, J. R., Jones, R., & Rothberg, S. J. (2001). Measurement of contact time in short duration sports ball impacts: An experimental method and correlation with the perceptions of elite golfers. *Sports Engineering*, 4, 191-203.
- Rose, A. (1942). The relative sensitivities of television pickup tubes, photographic film, and the human eye. *Proceedings of the Institute of Radio Engineers*, 30, 293-300.
- Runeson, S. (1983). On visual perception of dynamic events. In *Acta universitatis upsaliensis: Studia psychologica upsaliensia 9*.
- Runeson, S. (1995). Support for the cue-heuristic model is based on suboptimal observer performance: response to Gilden and Proffitt (1994). *Perception & Psychophysics*, 57, 1262-1273.
- Runeson, S., Juslin, P., & Olsson, H. (2000). Visual perception of dynamic properties: cue heuristics versus direct-perceptual competence. *Psychological Review*, 107, 525-555.
- Runeson, S., & Vedeler, D. (1993). The indispensability of precollision kinematics in the visual perception of relative mass. *Perception & Psychophysics*, 53, 617-632.
- Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to the rational model of categorization. *Psychological Review*, 117, 1144-1167.
- Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2009). A Bayesian framework for modeling intuitive dynamics. In *Proceedings of the 31st annual meeting of the cognitive science society*. Hillsdale, NJ: Erlbaum.
- Saxe, R., & Carey, S. (2006). The perception of causality in infancy. *Acta Psychologica*, 123(1-2), 144-65.
- Schlottmann, A., & Anderson, N. (1993). An information integration approach to phenomenal causality. *Memory and Cognition*, 21, 785-801.
- Schlottmann, A., & Shanks, D. R. (1992). Evidence for a distinction between judged and perceived causality. *Quarterly Journal of Experimental Psychology*, 44A, 321-34.
- Scholl, B. J., & Nakayama, K. (2002). Causal capture: Contextual effects on the perception of

- collision events. *Psychological Science*, *13*, 493-498.
- Scholl, B. J., & Nakayama, K. (2004). Illusory causal crescents: Misperceived spatial relations due to perceived causality. *Perception*, *33*(4), 455-469.
- Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, *4*(8), 299-309.
- Schwartz, O., Sejnowski, T. J., & Dayan, P. (2009). Perceptual organization in the tilt illusion. *Journal of Vision*, *9*, 1-20.
- Sekuler, R., Sekuler, A., & Lau, R. (1997). Sound alters visual motion perception. *Nature*, *385*(6614), 308.
- Shanon, B. (1976). Aristotelianism, Newtonianism and the physics of the layman. *Perception*, *5*(2), 241-243.
- Shepard, R. N. (1984). Ecological constraints on internal representations: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review*, *91*, 417-447.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317-1323.
- Shi, L., Griffiths, T. L., Feldman, N. H., & Sanborn, A. N. (2010). Exemplar models as a mechanism for performing Bayesian inference. *Psychological Bulletin and Review*, *17*, 443-464.
- Siegler, I. A., Bardy, B. G., & Warren, W. H. (2010). Passive vs. active control of rhythmic ball bouncing: The role of visual information. *Journal of Experimental Psychology. Human Perception and Performance*, *36*, 729-750.
- Smith, K. A., & Vul, E. (2012). Sources of uncertainty in intuitive physics. In *Proceedings of the 34th annual conference of the cognitive science society* (p. 995-1000).
- Stocker, A., & Simoncelli, E. (2006). Sensory adaptation within a Bayesian framework for perception. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems 18* (pp. 1289-1296). Cambridge, MA: MIT Press.
- Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human

- visual speed perception. *Nature Neuroscience*, 9, 578-585.
- Straube, B., & Chatterjee, A. (2010). Space and time in perceptual causality. *Frontiers in Human Neuroscience*, 4(April), 28.
- Stronge, W. J. (2000). *Impact mechanics*. Cambridge, UK: Cambridge University Press.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12, 49-100.
- Todd, J., & Warren, W. (1982). Visual perception of relative mass in dynamic events. *Perception*, 11, 325-335.
- Tremoulet, P. D., & Feldman, J. (2000). Perception of animacy from the motion of a single object. *Perception*, 29, 943-951.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14, 101-118.
- Wagenmakers, E. J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192-196.
- Warren, W. H., Kim, E. E., & Husney, R. (1987). The way the ball bounces: visual and auditory perception of elasticity and control of the bounce pass. *Perception*, 16(3), 309-336.
- Watanabe, K., & Shimojo, S. (2001). When sound affects vision: effects of auditory grouping on visual motion perception. *Psychological Science*, 12(2), 109-116.
- Weir, S. (1978). The perception of motion: Michotte revisited. *Perception*, 7, 247-260.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5, 598-604.
- White, P. A. (1988). Causal processing: Origins and development. *Psychological Bulletin*, 104(1), 36-52.
- White, P. A. (2006). The causal asymmetry. *Psychological Review*, 113(1), 132-147.
- White, P. A. (2009). Perception of forces exerted by objects in collision events. , 116, 580-601.
- White, P. A., & Milne, A. (1997). Phenomenal causality: Impressions of pulling in the visual perception of objects in motion. *American Journal of Psychology*, 110, 573-602.
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *TRENDS*

in Cognitive Sciences, 10, 301-308.

Zago, M., McIntyre, J., Senot, P., & Lacquaniti, F. (2009). Visuo-motor coordination and internal models for object interception. *Experimental Brain Research, 192*, 571-604.

Appendix A: Model Details

This appendix gives details of the direct perception and heuristic models of mass judgments. The details of the noisy Newtonian model of mass judgments and of causal impression are also given.

Direct perception model

The direct perception model assumes that observers perceive the true mass ratio and that this true mass ratio is then corrupted by zero-median noise (Runeson, 1983; A. Cohen, 2006). The mass ratio is assumed to lie on a logarithmic scale and Gaussian noise with mean zero and variance σ^2 is added to the logarithmic transform of the true mass ratio on each trial. We fit σ^2 to Experiments 1 and 2 of Todd and Warren (1982) finding that $\sigma = 0.634$ produced the maximum likelihood fit to the data.

Heuristic model

The salience function for this model has not been formally defined, as the model was intended to explain the data qualitatively rather than quantitatively. As in A. Cohen (2006), we can make quantitative predictions by choosing salience thresholds to match the data. The salience threshold for the ricochet is the angle at which a ricochet would be noticed. As any ricochet in our later experiments entails a 180° change of direction, we decided that a reversal of direction would always be salient. If both objects ricochet, then the final speed heuristic is used. For the final speed heuristic, the salience and default heuristic parameters are chosen to provide the best quantitative fit to the data of Experiments 1 and 2 of Todd and Warren (1982). We also include a guessing parameter that gives the probability that the model will make an unbiased guess instead of following the heuristics.

We fit the parameters to the data and found that the best-fitting parameters were a final speed ratio salience threshold of 1.58, a probability of using the final speed heuristic over the ricochet heuristic of 0.19, and the probability of unbiased guessing was 0.17.

Noisy Newton model

A prior distribution was placed on each of the unobserved variables. The prior on the coefficient of restitution gave equal weight to all possible values, while the prior distributions on the masses reflects the assumption that lower masses are more likely than higher masses.

$$e \sim \text{Uniform}(0, 1)$$

$$m_a \sim \text{Exponential}(1)$$

$$m_b \sim \text{Exponential}(1)$$

We model the observable variables, such as the initial and final velocities, in terms of physically motivated priors and psychophysically motivated observation noise. The priors we place on the true, noiseless initial velocities reflect the belief that slower velocities are more likely than faster velocities; this is a standard assumption in Bayesian models of velocity perception (Weiss et al., 2002; A. A. Stocker & Simoncelli, 2006). We only allowed initial velocities that result in the objects making contact under the assumption that object a begins on the left-hand side of the display and object b begins on the right-hand side of the display (we did not allow any $u_b > u_a$). Given the initial velocities, the prior on the noiseless final velocities are the values given by Newtonian mechanics, with the assumption that no external force is acting upon the system of two objects. These prior distributions have one parameter, σ^2 , that controls the strength of the prior expectation that objects tend to move slowly. We matched this parameter to human data in the mass judgment task using a value of $\sigma^2 = 4$, and kept it fixed for the causality judgment task.

$$\bar{u}_a \sim \text{Gaussian}(0, \sigma^2)$$

$$\bar{u}_b \sim \text{Gaussian}(0, \sigma^2)$$

$$\bar{v}_a = \frac{m_a u_a + m_b (u_b + e(u_b - u_a))}{m_a + m_b}$$

$$\bar{v}_b = \frac{m_b u_b + m_a (u_a + e(u_a - u_b))}{m_a + m_b}$$

$$\bar{g} = 0$$

$$\bar{t} = 0$$

Our observation model, linking true hidden velocities to observed velocities, follows the structure of threshold estimates of perceived velocity (Hick, 1950; Notterman & Page, 1957),

distance (Andrews & Miller, 1977), and time (Ekman, 1959). Qualitatively, it treats the true, noiseless value as the median of the observed distribution on values, with increasing noise levels for larger absolute magnitudes. We model this by adding a fixed amount of Gaussian noise to each velocity on a logarithmic scale. More precisely, let x be the variable, and \bar{x} be the true, noiseless value, which will become the median of the observed random variable

$$x = f^{-1}(f(\bar{x}) + \epsilon_x)$$

where $f(\bar{x}) = \text{sign}(\bar{x}) \log(w\bar{x} + 1)$, f^{-1} is the inverse of f , and $\epsilon \sim \text{Gaussian}(0, k_x^2)$. We used this formulation (e.g., A. A. Stocker & Simoncelli, 2006) instead of an unmodified log transformation to allow for both positive and negative velocities and to prevent a discontinuity at zero. For all variables we set $w = 0.15$, which results in a mild nonlinearity for the velocities, gaps, and time delays. We allowed the values of k to vary between variables, with $k_v = 0.1$ for the velocities, $k_g = 0.007$ for the gap, and $k_t = 0.003$ for the time delay. The values of k_v and w were set to match the human data in the mass judgment task and the values of k_g and k_t were set to match the spread of the gap and time delay judgments in the causality judgment task.

Qualitatively, then, we have constructed a probabilistic model with (1) no a priori bias on elasticities, (2) the expectation that extremely heavy or fast-moving objects are rare, (3) that collisions follow Newtonian mechanics (and non-collisions involve randomly generated velocities), and (4) that perceptual input is noisy (with more noise for faster velocities, larger gaps, or longer times). Bayesian inference in this model makes it possible for us to take the observed velocities and time delays and calculate the answer of a wide range of questions about all the other hidden variables. For example, we can calculate which object is heavier, and also the relative probability that a collision occurred at all.

We repeat our simulation many times for each trial seen by each participant. On a trial, we take the noiseless values from the display, find the probabilities of the various responses, and then make a response according to its probability. This approach is standard practice for probabilistic models of cognition (e.g., Anderson, 1991; Sanborn, Griffiths, & Navarro, 2010), making the assumption that people probability match to their internal belief distribution (Vulkan, 2000).

Another possibility is to sample a set of noisy observed variables based on the true observed variables for that trial, using the noise model above. Once these samples are acquired, the response made by the model is deterministic: the choice alternative with the maximum probability is always selected (e.g., A. Stocker & Simoncelli, 2006). We found the same qualitative results for our model implemented with the deterministic decision rule.

For mass judgments, the two hypotheses were $m_a > m_b$ and $m_b > m_a$. Given the noisy values of the velocities⁶, the probability of each hypothesis is determined and the hypothesis with higher posterior probability is selected. For causality judgments, the inputs to the model are the noisy values of the velocities, the gap, and the time delay. Given these values, the probability that the noisy Newton model produced these values is calculated. The alternative, that the values were generated randomly, is given a fixed probability p_r for all trials meaning that every trial has a uniform probability under this alternative model. The prior probabilities of the two hypotheses were assumed to be equal, so the probability that the values are generated randomly acts as a threshold to which the probability of the observations under the noisy Newton model is compared. If the probability that the values are produced from the noisy Newton model exceeded p_r , then the response is that it was a causal event, otherwise it is judged not a causal event. For the causality experiment, $p_r = 10^{-10}$.

Appendix B: Replication Details

This appendix gives the details of the two replications of the mass judgment experiments in Todd and Warren (1982).

Replication of Experiment 1 of Todd and Warren (1982)

Twenty four participants were recruited from a university community for this study. Four participants were discarded due to a computer error. Each participant was paid \$4 for less than one hour of participation. The eyes of the participants were situated approximately 44 cm away from the display.

⁶The perceived gap and time delay, even if noisy, do not impact this choice.

Participants were presented with movies of two white squares colliding with each other along one dimension. They were told that these squares were blocks sliding along an invisible smooth surface. Participants were instructed to press a key corresponding to whichever block they thought was heavier. The two white squares with 1 cm sides started outside the visible area of the screen and moved toward each other at their initial velocities until the edges of the two squares touched at the center of the screen. Following contact, the squares immediately moved away from each other at their final velocities. The trial ended automatically when the faster object reached the edge of the visible display, but participants could end the trial at any point by responding. No feedback was given to participants during the experiment.

A total of 252 trials were presented to each participant. There were twelve combinations of mass ratios and coefficients of restitution. One example of each combination was used presented to participants (order randomized for each participant) at the beginning of the experiment to acclimate them to the display. The data from these practice trials were not included in the analysis. The test trials consisted of 20 replications of each combination of mass ratio and coefficient of restitution with the order of presentation randomized for each participant. The mass ratios of the heavier to lighter object were set to be 1.25, 1.5, 2.0, or 3.0. On each trial, the heavier object was set to be the right or left object with equal probability. The coefficients of restitution used were 0.9, 0.5, and 0.1. The initial velocities of the left square ranged from 1.91 to 4.45 cm/s in steps of 0.13 cm/s. The initial velocity of the right square was determined by the initial velocity of the left square using the formula, $u_b = u_a - 6.35$ cm/s. The initial velocity of each trial was drawn uniformly from the set of initial velocities. Given these variables the final velocities of the two objects were uniquely determined.

Results are reported in the main text.

Replication of Experiment 2 of Todd and Warren (1982)

Twenty-two participants were recruited from a university community for this experiment. The equipment and experimental design of this experiment was identical to that of Experiment 1

with a few exceptions. The first change is that the initial velocity of the right square was fixed at zero. An additional level of coefficient of restitution was added $e = 1$ to the other three levels, resulting in sixteen combinations of mass ratio and coefficient of restitution. Sixteen practice trials began the experiment, one from each combination of mass ratio and coefficient of restitution, followed by 15 replications of each condition, giving a total of 256 trials.

Results are reported in the main text.

Appendix C: Explanation for Predictions of the Heuristic Model

The predictions of the heuristic model for Experiment 2 of Todd and Warren (1982) can be made clear by examining the ratio of final velocities of the two objects, as predicted by Newtonian mechanics. The final speed heuristic depends entirely on the magnitude of the ratio of the final velocities, while whether object a ricochets depends on the sign of the final velocity ratio. Because $u_b = 0$, the relationship between the mass ratio and the ratio of the final velocities is a generalization of the relationship given in Gilden (1991)

$$\frac{v_b}{v_a} = \frac{\frac{m_a}{m_b}(1+e)}{\frac{m_a}{m_b} - e}.$$

Figure 14 shows the ratio of the final velocities by mass ratio for different values of e . Each velocity ratio has a discontinuity marked by a vertical line in the figure, and this discontinuity is where $m_a/m_b = e$. On the right-hand side of the discontinuity, $|v_b| > |v_a|$, indicating that $m_a > m_b$ by the final speed heuristic. As there is a lack of a ricochet on the right-hand side of the discontinuity, the heuristic model will always choose $m_a > m_b$ if the final speed exceeds the salience threshold. On the left-hand side of the discontinuity, the trials in our experiment all fall into the region in which $|v_b| \geq |v_a|$, but a ricochets. The heuristics contradict each other in this region and can produce any response proportion for an individual mass ratio by manipulating the salience function and proportion each heuristic is the default in the population of participants.

The ordering of the discontinuities allow for the model to produce the ordering of subjective points of equality of Experiment 2 of Todd and Warren (1982): a point of subjective inequality that is more biased towards $m_a > m_b$ as the value of e decreases. However, there is a mismatch

with human data because on the right-hand side of the discontinuity, as m_a/m_b increases, the ratio v_b/v_a decreases. This relationship means that the heuristic model predicts that as m_a/m_b becomes larger, the probability of choosing $m_a > m_b$ will remain the same or decrease. This holds even if more general monotonic salience functions are used instead of salience thresholds. This prediction does not match the data for the condition $e = 0.1$, and as it is due to a structural property of the heuristic model; the discrepancy cannot be removed by fitting parameters to individual participants and averaging.




Figure 1. Experimental designs for the mass and collision judgment tasks. A) Stages of a mass judgment movie. In the first row, two objects have initial velocities u_a and u_b and masses m_a and m_b . The second row shows the collision. The final velocities v_a and v_b are shown in the last row. B) Stages of a causality judgment movie. In the first row, the gray square is moving with initial velocity u_a and the black square is stationary. Both squares have the same mass m . The second and third rows show the gap g between squares after the first square stops and the time delay t between the first square stopping and the second square starting to move. The final row shows the gray square stationary and the black square moving with final velocity v_b .




Figure 2. Graphical model showing the dependency between variables for the noisy Newton model of mass judgments. Observable variables are shaded gray. The observed variables, initial velocities u_a and u_b , final velocities v_a and v_b , each depend on an independent source of noise ϵ . The final velocities v_a and v_b also depend on the object masses m_a and m_b and coefficient of restitution e . The dashed double-headed line between u_a and u_b restricts the noiseless values of these variables to values for which the objects will make contact: $u_b < u_a$.

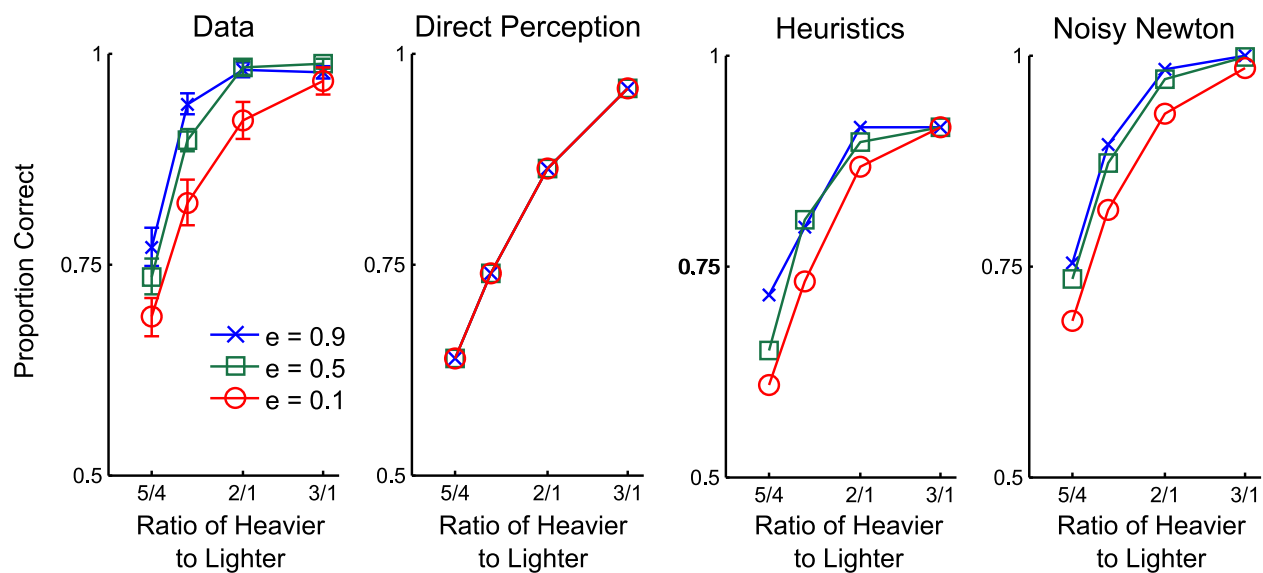


Figure 3. Data and modeling results for the replication of Experiment 1 of Todd and Warren (1982). The horizontal axis is the mass ratios of the collisions shown to participants and the vertical axis is mean response accuracy. The separate lines correspond to different coefficients of restitution in the collisions. Error bars are plus or minus one standard error of the mean.

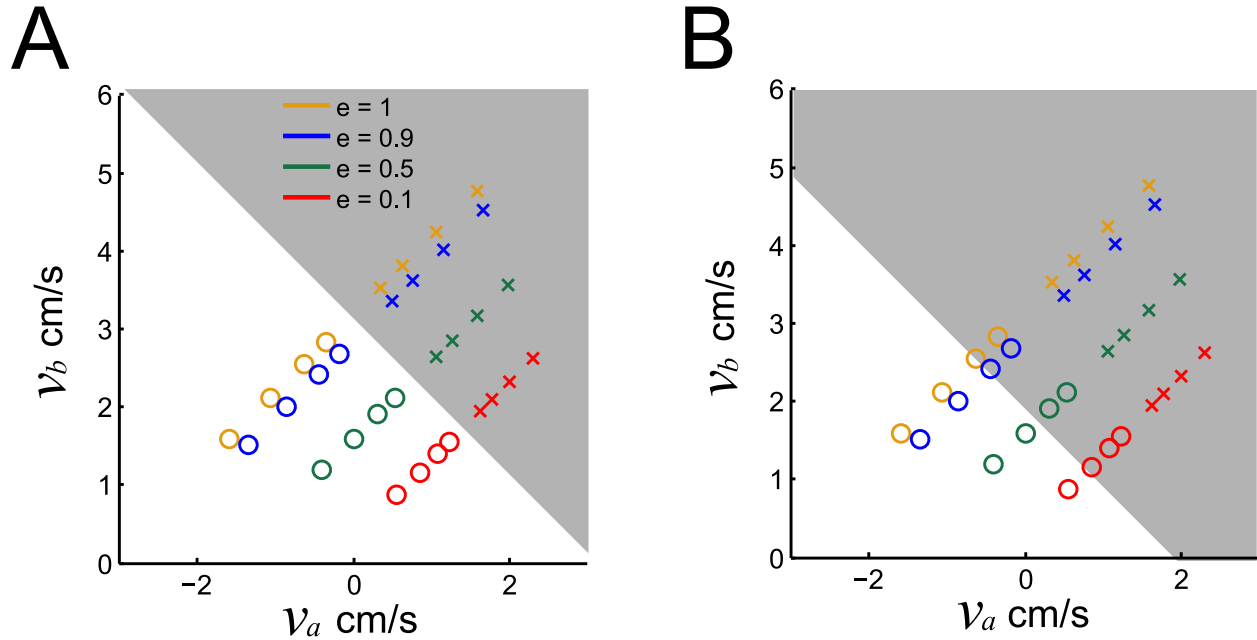


Figure 4. Explanation for the predictions of the noisy Newton model in the mass judgment task.

A) Decision boundary for a mass judgment experiment based on perfect perception of the initial velocities. The x's in each plot are the conditions in which $m_a > m_b$ and the o's are the conditions in which $m_a < m_b$. The markers are colored by the coefficient of restitution condition and the markers that are closer to the decision boundary have lower mass ratio. The gray area shows where $m_a > m_b$ and the decision boundary is the edge between the gray and white areas. Note that a lower coefficient of restitution brings the markers closer to the boundary. B) Decision boundary for a mass judgment experiment based on a decrease in the inferred sum of initial velocities.

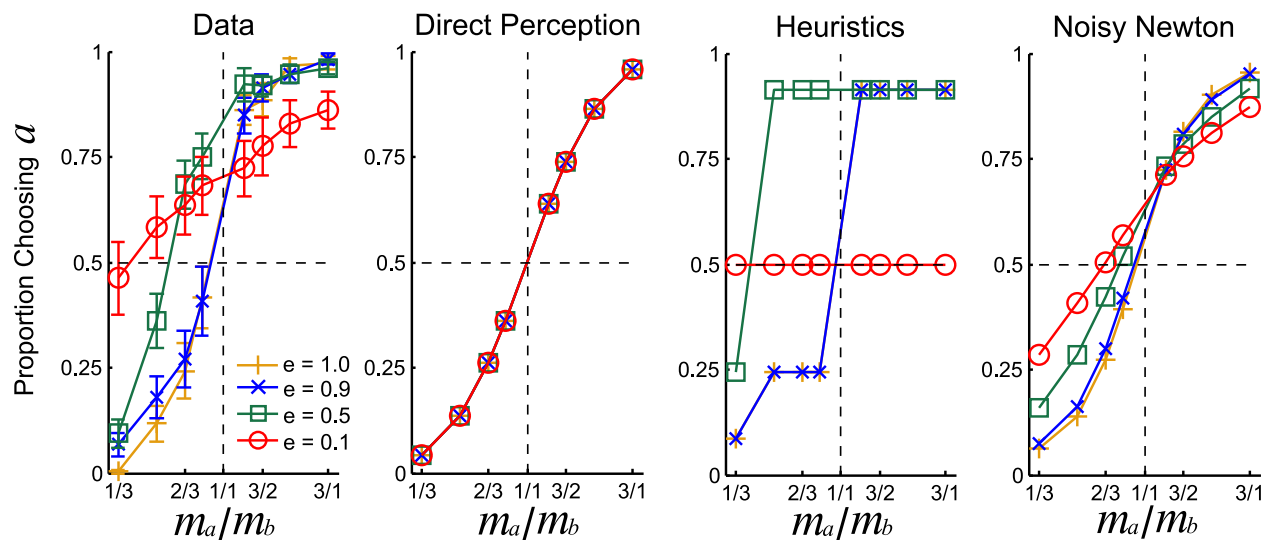


Figure 5. Data and modeling results for the replication of Experiment 2 of Todd and Warren (1982). The horizontal axis is the mass ratio of the initially moving object to the initially stationary object of the collisions shown to participants. The vertical axis is the proportion of trials on which participants chose $m_a > m_b$. The point of subjective equality is where the data lines cross the horizontal dotted line. A bias is shown if the data lines do not cross the horizontal dotted line at the vertical dotted line.

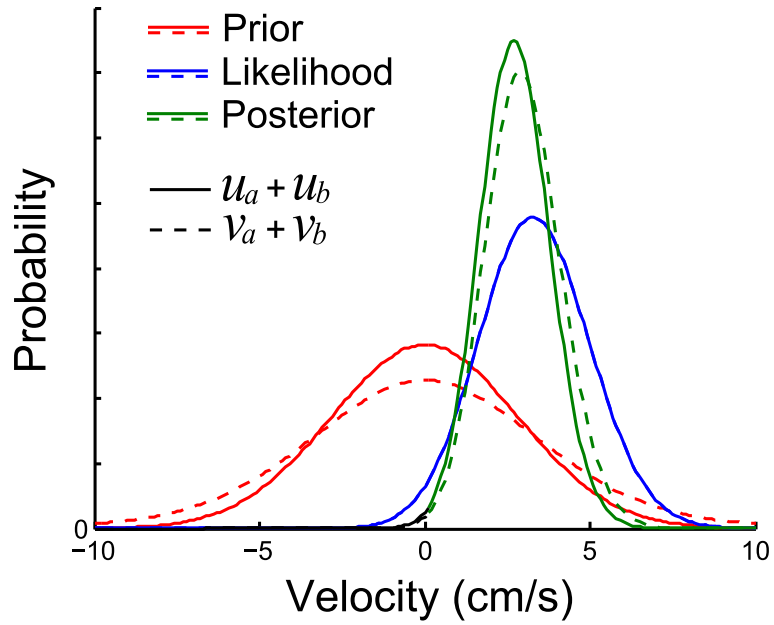


Figure 6. Comparison of the prior, likelihood, and posterior densities over the sum of the initial velocities, $u_a + u_b$, to the same densities over the sum of the final velocities, $v_a + v_b$. Line color indicates whether the density function is a prior, likelihood, or posterior density. Line style indicates whether a density function is over $u_a + u_b$ or over $v_a + v_b$.

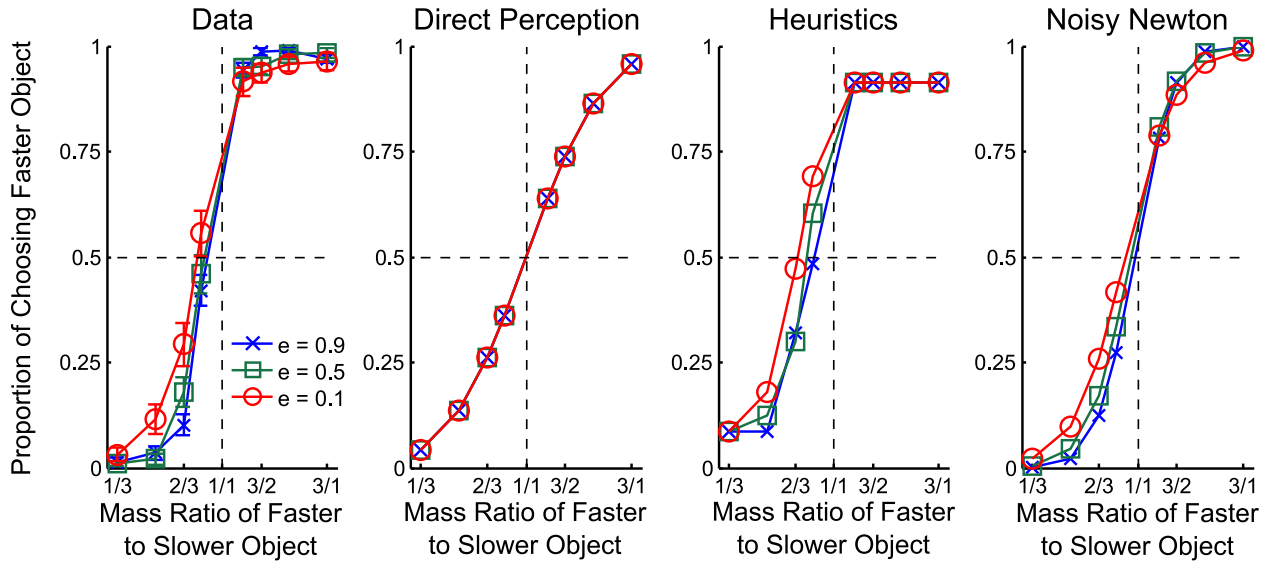


Figure 7. Data and modeling results for the replication of Experiment 1 of Todd and Warren (1982) plotted to show bias. The horizontal axis is the ratio of the mass of the initially faster moving object to the mass of the initially slower moving object. The vertical axis shows proportion of responses participants made indicating that the faster object had greater mass. Where the data lines cross the horizontal dotted lines gives the point of subjective equality. The separate lines correspond to different coefficients of restitution in the collisions. Error bars are plus or minus one standard error of the mean.

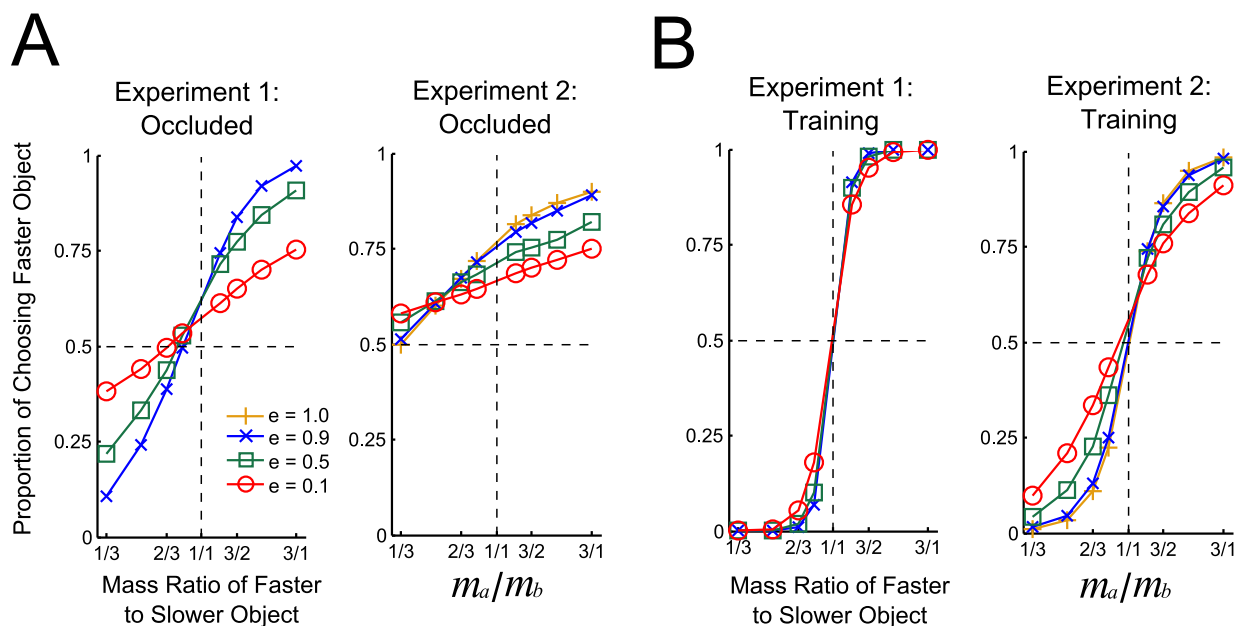


Figure 8. Model predictions of the noisy Newton model for occlusion or a prior distribution based on training in the replications of Experiments 1 and 2 of Todd and Warren (1982). The horizontal axis is the ratio of the mass of the initially faster moving object to the mass of the initially slower moving object of the collisions shown to participants for all plots. The vertical axis is the proportion of responses picking the initially faster moving object. Where the data lines cross the horizontal dotted lines gives the point of subjective equality. The separate lines correspond to different coefficient of restitution conditions. A) Model predictions for occluded initial velocities. B) Model predictions with a prior that matches a set of training trials.

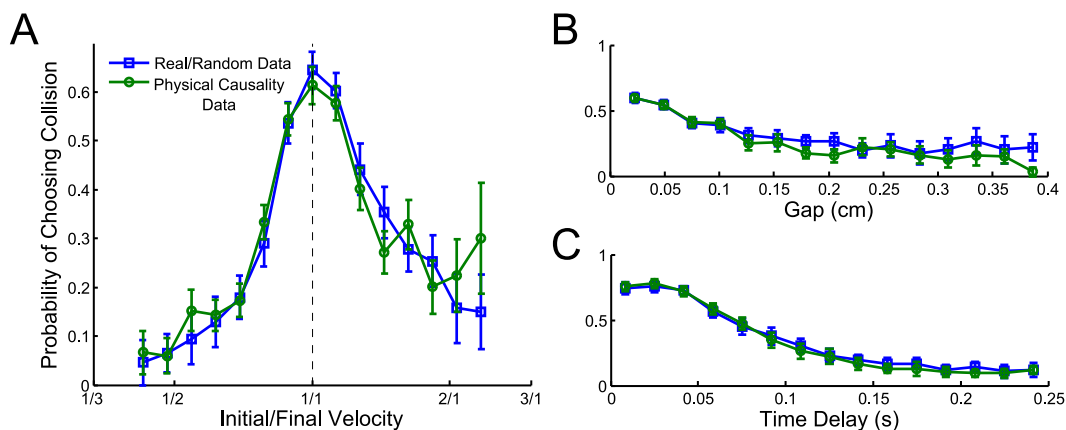


Figure 9. Results of the causality judgment task. A) Human data from the Real or Random and Physical Causality conditions for final velocity. The ratio of the initial velocity of Object A over the final velocity of Object B lies along the horizontal axis. The vertical axis is the probability of choosing that the collision was causal or a real collision with data aggregated over participants. The blue squares are the binned data from the Real or Random condition and the green circles are the binned data from the Physical Causality condition. Error bars are plus or minus one standard error of the mean over subjects. B) Human data from the Real or Random and Physical Causality conditions for gap. C) Human data from the Real or Random and Physical Causality conditions for time delay.




Figure 10. Graphical model showing the dependency between variables for the noisy Newton model for causality judgments. Observable variables are shaded gray. The observed variables, initial velocities u_a and u_b , final velocities v_a and v_b , gap g , and time delay t , each depend on an independent source of noise ϵ . The final velocities v_a and v_b also depend on the object masses m_a and m_b and coefficient of restitution e . The dashed double-headed line between u_a and u_b restricts the noiseless values of these variables to values for which the objects will make contact: $u_b < u_a$. The dashed double-headed line between m_a and m_b indicates an extra constraint (when necessary) between the two masses: $m_a = m_b$.

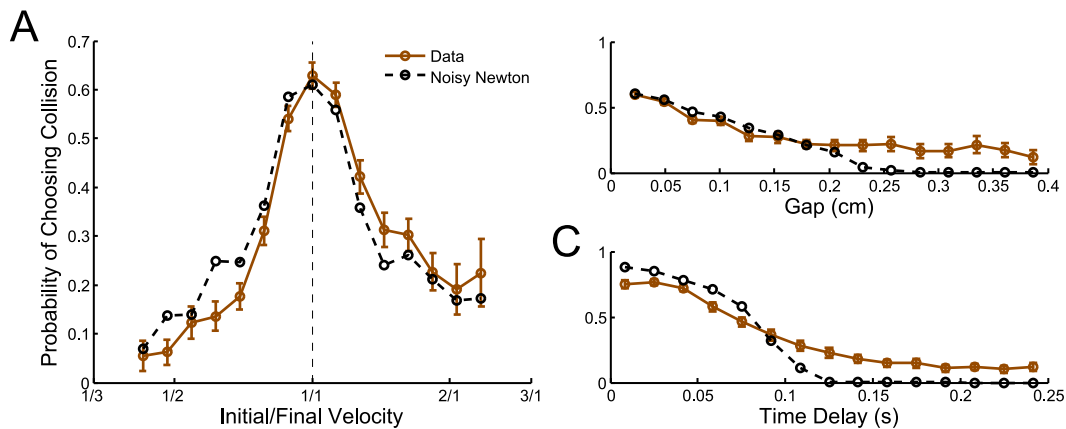


Figure 11. Results and model fit of the first causality judgment experiment. A) Human data and model predictions for final velocity. The ratio of the initial velocity of Object A over the final velocity of Object B lies along the horizontal axis. The vertical axis is the probability of choosing causality or a real collision over the alternative with data aggregated over participants and conditions. The brown solid line is the binned data; the black dashed line is the model predictions. Error bars are plus or minus one standard error of the mean over subjects. B) Human data and model predictions for gap. C) Human data and model predictions for time delay.

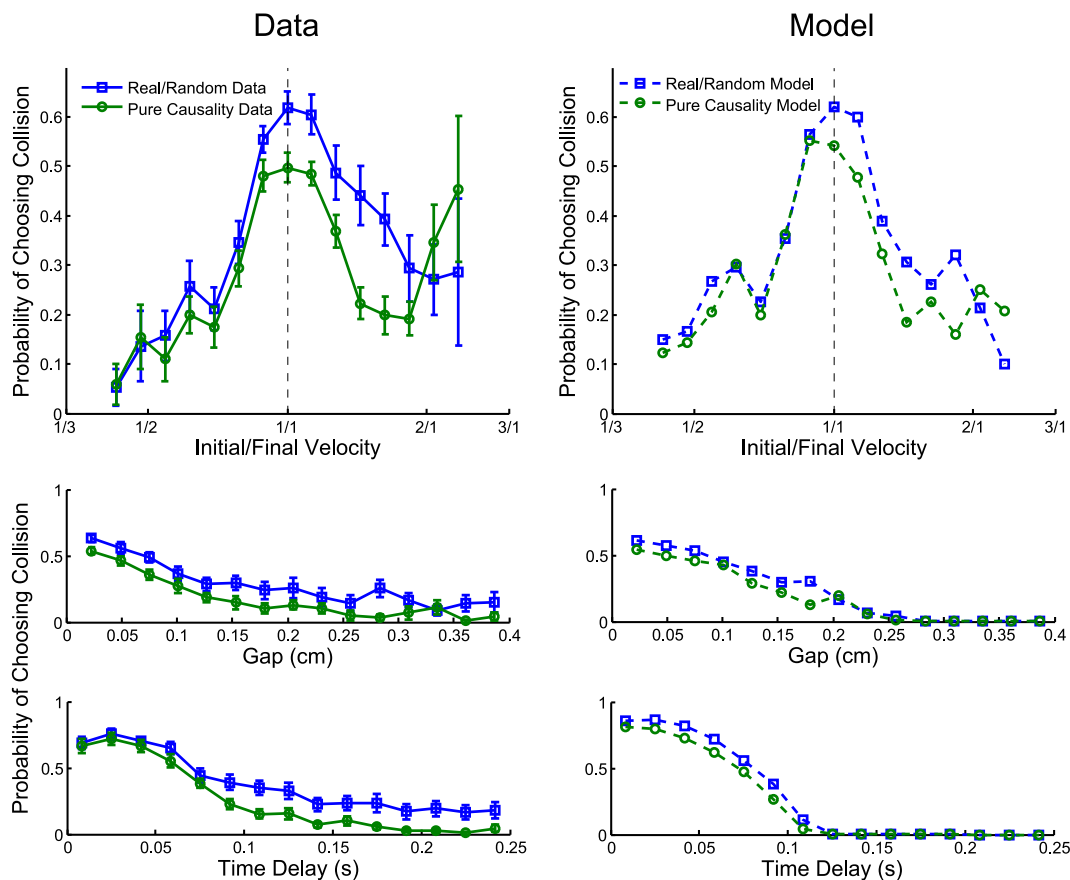


Figure 12. Results and model fit of the causality judgment task in the second experiment. Both columns present the Real or Random and Pure Causality conditions for final velocity, gap, and time delay. The vertical axis is the probability of choosing that the collision was causal or a real collision with data aggregated over participants. The blue squares are the binned data from the Real or Random condition and the green circles are the binned data from the Pure Causality condition. Error bars are plus or minus one standard error of the mean over subjects. The data column presents the human data and the model column presents the noisy Newton model predictions.

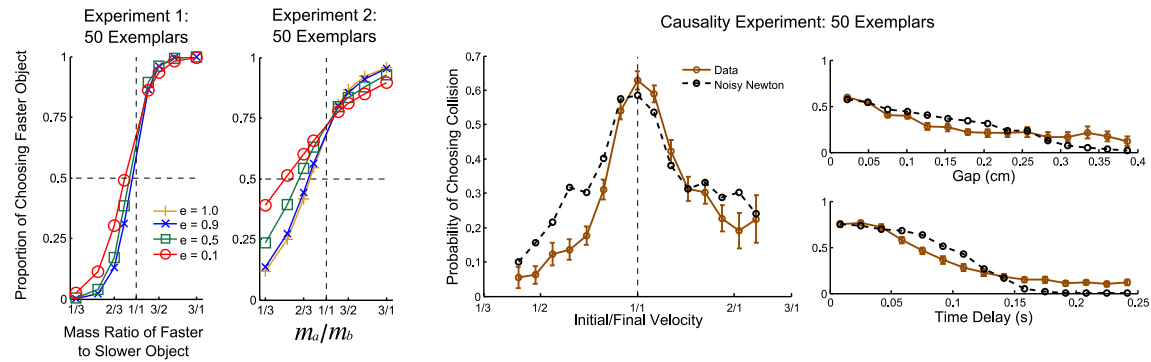


Figure 13. Model predictions of the noisy Newton model using a prior distribution that consists of fifty noisy samples from the hypothesized environment. For the mass judgments, the predictions are from the replications of Experiments 1 and 2 from Todd and Warren (1982). The horizontal axis is the ratio of the mass of the initially faster moving object to the mass of the initially slower moving object of the collisions shown to participants for both plots. The vertical axis is the proportion of responses picking the initially faster moving object. Where the data lines cross the horizontal dotted lines gives the point of subjective equality. The separate lines correspond to different coefficient of restitution conditions. The causality data are from our first causality experiment. For these data, the vertical axis is the probability of choosing that the collision was causal or a real collision with data aggregated over participants and conditions. The brown solid line is the binned data; the black dashed line is the model predictions. Error bars are plus or minus one standard error of the mean over subjects.




Figure 14. Final velocity ratio by mass ratio for the heuristic model. Each colored line indicates a different coefficient of restitution, with the vertical lines marking the discontinuities in the functions. Dashed or dotted lines separate regions where the final speed and ricochet heuristics combine in different ways. The gray areas indicate where the final speed heuristic predicts that $m_a > m_b$.